Volume 13, Number 3

September 2017

IJCB

INTERNATIONAL JOURNAL OF CENTRAL BANKING

Special Issue: Challenges to Financial Stability in a Low Interest Rate World

IJCB
INTERNATIONAL JOURNAL OF CENTRAL BANKING
Volume 13, Number 3
September 2017

IJCB International Journal of Central Banking

Volume 13, Number 3 September 2017

Special Issue on Challenges to Financial Stability in a Low Interest Rate World

The contents of this journal, together with additional materials provided by article authors, are available without charge at ***www.ijcb.org****.*

Copyright © 2017 by the Association of the International Journal of Central Banking. All rights reserved. Brief excerpts may be reproduced or translated provided the source is cited. Consult www.ijcb.org for further information.

The *International Journal of Central Banking* is published quarterly (ISSN: 1815-4654). Online access to the publication is available free of charge at **www.ijcb.org**. Individual print subscriptions are available. Orders may be placed by phone (001 415 974 2035), via fax (001 415 974 2168), or by e-mail (editor@ijcb.org).

Renewals, claims, address changes, and requests for permission to reprint material from this journal should be addressed to:

International Journal of Central Banking
Economic Research Department
Federal Reserve Bank of San Francisco
101 Market Street
San Francisco, CA 94105
USA

Phone: 001 415 974 2035
Fax: 001 415 974 2168
E-mail: editor@ijcb.org

The views expressed in this journal do not necessarily represent the views of the Association of the International Journal of Central Banking or any of its members.

ISSN: 1815-4654

International Journal of Central Banking

Board of Directors

Chairman

Claudio Borio, *Bank for International Settlements*

Board Members

Q. Farooq Akram, *Norges Bank*
Abdulaziz Al-Furaih, *Saudi Arabian Monetary Agency*
David E. Altig, *Federal Reserve Bank of Atlanta*
Carlos Hamilton Vasconcelos Araujo, *Central Bank of Brazil*
Jan Marc Berk, *The Nederlandsche Bank*
Mohamed Tahar Bouhouche, *Bank of Algeria*
Lillian Cheung, *Hong Kong Monetary Authority*
Laurent Clerc, *Bank of France*
Francisco G. Dakila Jr., *Central Bank of the Philippines*
Mary Daly, *Federal Reserve Bank of San Francisco*
Michael Dotsey, *Federal Reserve Bank of Philadelphia*
Gabriel Fagan, *Central Bank of Ireland*
Jiang Feng, *People's Bank of China*
Manuel Ramos Francia, *Bank of Mexico*
Jeffrey C. Fuhrer, *Federal Reserve Bank of Boston*
Kamil Galuscak, *Czech National Bank*
Niels Lynggård Hansen, *Danmarks Nationalbank*
Philipp Hartmann, *European Central Bank*
Elena Iorga, *National Bank of Romania*
Seung-Cheol Jeon, *Bank of Korea*
Juan F. Jimeno, *Bank of Spain*
George A. Kahn, *Federal Reserve Bank of Kansas City*
Sujit Kapadia, *Bank of England*
Ali Hakan Kara, *Central Bank of Turkey*
Christopher Kent, *Reserve Bank of Australia*
Evan Koenig, *Federal Reserve Bank of Dallas*
Thomas Laubach, *Federal Reserve Board*
Ana Christina Leal, *Bank of Portugal*
Carlos Lenz, *Swiss National Bank*
Jesper Lindé, *Sveriges Riksbank*
Choy Keen Meng, *Monetary Authority of Singapore*
John McDermott, *Reserve Bank of New Zealand*
Emanuel Moench, *Deutsche Bundesbank*
Alberto Naudon, *Central Bank of Chile*
Edward Offenbacher, *Bank of Israel*
Fabio Panetta, *Bank of Italy*
Michael D. Patra, *Reserve Bank of India*
Thórarinn G. Pétursson, *Central Bank of Iceland*
Ivan Ribnikar, *Bank of Slovenia*
Lawrence Schembri, *Bank of Canada*
Mark Schweitzer, *Federal Reserve Bank of Cleveland*
Shigenori Shiratsuka, *Bank of Japan*
Jan Smets, *National Bank of Belgium*
Daniel Sullivan, *Federal Reserve Bank of Chicago*
Juha Tarkka, *Bank of Finland*
George Tavlas, *Bank of Greece*
Joseph Tracy, *Federal Reserve Bank of New York*
Dobieslaw Tymoczko, *National Bank of Poland*
Hernando Vargas Herrera, *Banco de la República*
Christopher Waller, *Federal Reserve Bank of St. Louis*
John Weinberg, *Federal Reserve Bank of Richmond*
Mark Wright, *Federal Reserve Bank of Minneapolis*
Ksenia Yudaeva, *Central Bank of Russian Federation*

Editorial Board

Managing Editor

Loretta Mester
Federal Reserve Bank of Cleveland

Co-editors

Tobias Adrian
International Monetary Fund

Boragan Aruoba
University of Maryland

Pierpaolo Benigno
LUISS Guido Carli

Harrison Hong
Columbia University

Luc Laeven
European Central Bank

Lucrezia Reichlin
London Business School

Rafael Repullo
CEMFI

Barbara Rossi
Universitat Pompeu Fabra

Associate Editors

Patrick Bolton
Columbia University

Michael D. Bordo
Rutgers University

Mark Carey
Federal Reserve Board

Pierre Collin-Dufresne
Columbia University

Guy Debelle
Reserve Bank of Australia

Douglas W. Diamond
University of Chicago Graduate School of Business

Francis Diebold
University of Pennsylvania

Michael Dotsey
Federal Reserve Bank of Philadelphia

Darrell Duffie
Stanford University

Jordi Galí
Centre de Recerca en Economia Internacional (CREI)

Marvin Goodfriend
Carnegie Mellon University

Michael B. Gordy
Federal Reserve Board

Luigi Guiso
European University Institute

Andrew G. Haldane
Bank of England

Takatoshi Ito
Columbia University

David Lando
Copenhagen Business School

Philip Lane
Trinity College Dublin

Francesco Lippi
University of Sassari

Carmen M. Reinhart
Harvard Kennedy School

Eli M. Remolona
Bank for International Settlements

Hélène Rey
London Business School

Jean-Charles Rochet
University of Zurich

Andrew K. Rose
University of California, Berkeley

Klaus Schmidt-Hebbel
Organisation for Economic Co-operation and Development (OECD)

Lars E.O. Svensson
Stockholm School of Economics

Jürgen von Hagen
University of Bonn

Ernst-Ludwig von Thadden
University of Mannheim

Tsutomu Watanabe
University of Tokyo

Advisory Board

Franklin Allen
The Wharton School of the University of Pennsylvania

Charles Goodhart
London School of Economics

Hyun Shin
Bank for International Settlements

Kazuo Ueda
University of Tokyo

Carl E. Walsh
University of California

Michael Woodford
Columbia University

John Taylor
Stanford University

Are Low Real Interest Rates Here to Stay?*

Lukasz Rachel and Thomas D. Smith
Bank of England

Long-term real interest rates across the world are low, having fallen by about 450 basis points (bps) over the past thirty years. To understand whether low real rates are here to stay, we need to understand what has caused the decline. The co-movement in rates across both advanced and emerging economies suggests a common driver: the global neutral real rate may have fallen. In this paper we attempt to identify which secular trends could have driven such a fall. Although there is huge uncertainty, under plausible assumptions we think we can account for around 400 bps of the 450 bps fall. Our quantitative analysis highlights slowing global growth expectations as one force that may have pushed down on real rates recently, but shifts in saving and investment preferences appear more important in explaining the long-term decline. We think the global saving schedule has shifted out in recent decades due to demographic forces, higher inequality, and, to a lesser extent, the glut of precautionary saving by emerging markets. Meanwhile, desired levels of investment have fallen as a result of the falling relative price of capital, lower public investment, and an increase in the spread between risk-free and actual interest rates. Looking ahead, in the absence of sustained changes in policy, most of these forces look set to persist, and some may even build further. This suggests that the global neutral rate may remain low and perhaps settle at around 1 percent in the medium to

*This paper is based on material from Bank of England Staff Working Paper No. 571 (Rachel and Smith 2015), which contains additional detail, charts, and analysis. The views expressed in this paper are those of the authors and not necessarily those of the Bank of England or its committees. We would like to thank Scott Simmons for his research assistance during the project and colleagues in the Bank of England's International Directorate for helpful comments. All remaining errors and omissions are our own. Author e-mails: lukasz.rachel@bankofengland.co.uk, thomas.smith@bankofengland.co.uk.

long run. If true, this will have widespread implications for policymakers—not least in how to manage the business cycle if monetary policy is frequently constrained by the zero lower bound.

JEL Codes: E02, E10, E20, E40, E50, E60, F00, F41, F42, F47, J11, O30, O40.

1. Introduction

The downward trend in long-term risk-free interest rates is not a new phenomenon. Alan Greenspan famously highlighted the decline in long-term U.S. bond yields in 2005, which occurred despite the Federal Reserve tightening policy. As Greenspan noted, while the downward trend in yields was clear, the explanations for the fall were not—it was a "conundrum." This conundrum is not unique to the United States either. Since the 1980s, market measures of long-term risk-free real interest rates have declined by around 450 bps across both emerging and developed economies (figure 1). Although there is a lot of variation across countries, the presence of a discernible common trend suggests that global factors are at work.

The decline in global real interest rates has largely occurred against a backdrop of low and stable inflation with little sign of demand overheating. This suggests that the sustained fall in long-term market rates is symptomatic of a fall in the global neutral rate. The global neutral rate is an important policy variable, as it acts as an anchor for a country's equilibrium real rate in the long run.[1]

The purpose of this paper is to contribute to the wider debate about why the global neutral rate may have fallen and what may happen going forward (Summers 2014a, 2014b). Our main contributions are twofold: first, we assemble a rich collection of global data to analyze the main secular trends that could be driving the global neutral real rate. Second, we develop a simple accounting framework to quantify the relative importance of these trends in a coherent way. We then use these insights to explain the fall in the global neutral

[1] For a great overview of the importance of the equilibrium rate for policy, see Fischer (2016).

Figure 1. Comparison of Real Interest Rates

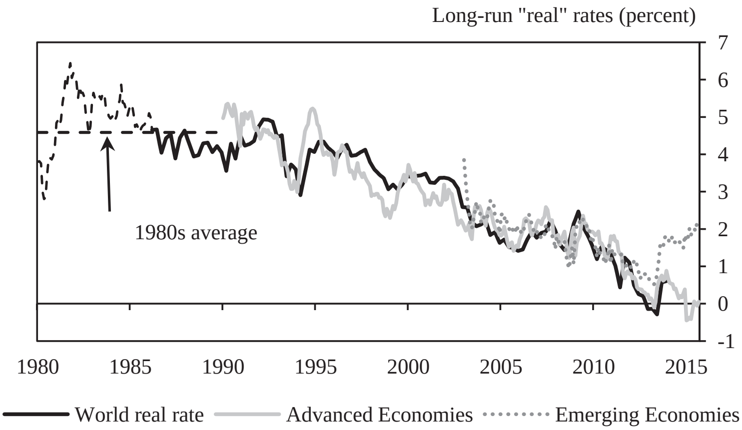

Notes: The "world" real rate (solid black line) is taken from King and Low (2014) and shows the average ten-year yield of inflation-linked bonds in the G7 countries (excluding Italy) over the period 1985–2013. This line has been extended back to the start of the 1980s (dashed black line) using a simple regression linking it to movements in UK ten-year nominal yields and RPI inflation. The solid and dotted grey lines show simpler measures of real rates for different country groups, calculated as the nominal yield on ten-year sovereign bonds minus one-year-ahead inflation expectations from Consensus Economics. Figures have been GDP-weighted together for twenty advanced economies (solid grey) and seventeen emerging markets (dotted grey).

rate in the past, and offer a prediction of how the neutral rate could evolve in the future.

The global neutral rate is largely determined by expectations of global trend growth and other factors shaping preferences for desired saving and investment. We analyze each in turn. First we use a modified growth-accounting framework to analyze the various secular trends that could be affecting global growth. Then we use a simple saving-investment framework to analyze global shifts in desired saving and investment to analyze how changes in preferences could have affected the neutral rate.

Although changes in global trend growth are probably the most commonly cited driver of changes in real interest rates, we find it difficult to account for much (if any) of the pre-crisis fall in global real rates by just appealing to past changes in growth—global growth was fairly steady in the pre-crisis decades.[2] However, the financial crisis does appear to have triggered a wider reassessment of growth prospects, and lower expectations of future growth seem to be playing a role in driving the more recent decline in real rates. Our analysis suggests that slower global labor supply growth (due to demographic forces) and headwinds at the technological frontier (such as a plateau in educational attainment) may cause global growth to slow by up to 1 percentage point (pp) over the next decade. We think expectations of this decline could account for about 100 bps of the fall in real rates seen recently.

Shifts in the balance of desired saving and investment appear quantitatively even more important than changes in growth expectations. Our analysis suggests that the desired saving schedule has shifted out materially due to demographic forces (90 bps of the fall in real rates), higher inequality within countries (45 bps), and a preference shift towards higher saving by emerging market governments following the Asian crisis (25 bps). If this had been the whole story, we would have expected to see a steady rise in actual saving rates globally. But global saving and investment ratios have been remarkably stable over the past thirty years, suggesting that desired investment levels must have also fallen. We pin this decline in desired investment on a fall in the relative price of capital goods (accounting for 50 bps of the fall in real rates) and a preference shift away from public investment projects (20 bps). Also, we note that the rate of return on capital has not fallen by as much as risk-free rates. The rising spread between these two rates has further reduced desired investment and risk-free rates (by 70 bps). Together these effects can account for 300 bps of the fall in global real rates.

When combined, lower expectations for trend growth and shifts in desired saving and investment can account for about 400 bps of the 450 bps decline in the global long-term neutral rate since the late 1980s. Even more difficult than accounting for the past is predicting what might happen from here. Absent major policy changes, our analysis suggests that many of these secular trends look likely to persist and some may even build further. If so, the global neutral

[2]This is consistent with the finding that historically there is only a weak relationship between realized GDP growth and the real interest rate—see Hamilton et al. (2015).

rate may stay low, settling at around 1 percent over the medium to long run.

The rest of this paper is structured as follows. Section 2 starts by describing the concept of the global neutral rate as a long-run anchor for country-specific equilibrium policy rates and the data we use to measure movements in the neutral rate over time. Section 3 discusses the role of economic growth in driving real rates over the past and future. Section 4 analyzes the role of shifts in desired saving and investment as a driver of real rates. Finally, section 5 concludes.

2. Concept and Data

In order to analyze changes in the (unobservable) global neutral rate over time, we need to define what the global neutral rate is and establish an empirical basis for how it has moved.

We define the global neutral rate as the rate to which country-specific equilibrium rates will converge in the long term, absent distortions and shocks. Put differently, the global neutral rate acts as an anchor for equilibrium real rates in open economies.[3] In reality, plenty of distortions and shocks will drive a wedge between country-specific real rates and this long-term anchor. These can be divided into global factors and country-specific factors.[4] Among the global factors, it is useful to distinguish between persistent headwinds that can take several years or even decades to subside (such as the global deleveraging process under way since the crisis) and short-run global cyclical factors (such as global credit conditions or levels of confidence).[5] Among the country-specific factors, a country's cyclical position could drive a country's real rates temporarily higher or lower than the global level. Additionally, a country's structural characteristics—such as its demographic structure, its trend

[3]See Mendes (2014) for an intuitive exposition of the model where a small open economy is a price taker in the global market.

[4]See Mishkin (1984) for evidence of how country-specific risk factors have prevented equalization of real returns across countries.

[5]The evolution of the global neutral real rate since the global financial crisis is particularly uncertain because it is too early to tell how much of the most recent decline is cyclical, and hence will prove temporary, or persistent, but not reflective of the long-term neutral rate.

rate of productivity growth, or the quality of its institutions—may drive individual country rates persistently above or below global R*. An individual country's monetary policy stance can also temporarily drive a wedge between a country's equilibrium real rate and the actual (real) policy rate.[6]

In this paper we use ten-year government bond yields, adjusted for inflation, as an indicator for how the global neutral rate has moved over time. These market measures are subject to the same shocks and distortions as those for country-specific equilibrium rates (highlighted above) as well as others such as the impact monetary policy regime shifts can have on term premia and distortions specific to government bond markets. To account for such distortions, we make a number of adjustments in our analysis.

To smooth out the effect from country-specific factors, we focus our analysis on global measures of real rates aggregated across countries. To sidestep cyclical issues, we focus our analysis on very low frequency movements in global data. For example, in section 3 we consider five-year averages of the global growth rate, and in section 4 we focus primarily on decade averages of the real interest rate. To avoid the impact that monetary policy regime shifts may have had on term premia, we focus our analysis on explaining changes in real rates since the average of the 1980s, rather than the 1980s peak. The average of the 1980s also corresponds to the point where real rates hovered in the early 1990s, which is after many of the largest monetary policy regime shifts had occurred (see International Monetary Fund 2014b).[7]

[6]For a model-based estimate of the equilibrium real interest rate which is similar in spirit to what we focus on in the present paper, see Bomfim (1997). For an example of how the equilibrium rate features in policy briefings and deliberations, see the Federal Open Market Committee briefing documents available at https://www.federalreserve.gov/monetarypolicy/files/fomc20010515bluebook 20010510.pdf.

[7]In principle, long-term forward rates would be a better market measure to use in this analysis than long-term spot rates, as forward rates exclude some of the distortions associated with term premia. A lack of historical data on forward rates across countries hinders such analysis at the global level, but King and Low (2014) show that for countries where data are available (such as the United Kingdom), forward rates have tended to co-move with spot rates over the past thirty years.

Changes to the structure of government bond markets could also have affected long-term yields over time. Over the past thirty years there have been significant shifts in both the demand and supply of government bonds. In the decades before the crisis, the supply of high-quality government debt increased steadily. This rise in supply was partly met by increased demand for safe assets by emerging market governments, but in net terms the supply of safe assets increased in the decades before the crisis, potentially putting upward pressure on government bond yields and counteracting some of the fall in long rates observed in the data. However, since the crisis that picture has changed markedly. Deteriorating fiscal positions and ratings downgrades have seen the supply of high-quality debt fall sharply, while demand for safe assets has risen as a result of permanent regulatory changes and cyclical central bank action. This suggest that since the crisis it may well be true that the net supply of safe assets has deteriorated (Caballero and Farhi 2013), and this may have contributed to the fall in actual real interest rates seen recently. However, we still need to appeal to other factors to explain the steady decline in real rates before the crisis and arguably some of the fall since—as highlighted in sections 3 and 4 below.

Another potential issue with market measures of global real rates based on long-term bond yields is that they are affected by low short-term interest rates. Some have even argued that "global rates are low because monetary policy is loose." If this were the case, low rates would not pose a policy dilemma. The solution would be trivial and the downward trend in long rates would simply reverse when central banks tightened policy. But as Bernanke (2015a) and Broadbent (2014) have pointed out, this view of the world is unlikely to be correct, as the decline in actual real interest rates has occurred against a backdrop of contained inflation with little sign of exuberant demand growth. Indeed, global growth and inflation have, if anything, disappointed in the most recent recovery, despite policy rates being historically low. This suggests that observable interest rates have merely followed their unobservable "equilibrium" counterparts—if policy had been tighter, inflation would have been lower and demand would have been too weak to deliver full employment. Ever looser monetary policy is not the cause but the consequence of the fall in long-term rates.

This intuition has been formalized in econometric models, which aim to extract measures of equilibrium interest rate from observed data. Laubach and Williams (2003) perform this sort of exercise for the United States and find that U.S. R* has declined by around 450 bps since the 1960s, and by around 300 bps since the 1980s. The authors suggest that secular trends related to changes in trend growth and shifts in saving and investment preferences are responsible for this decline—not monetary policy. We use a similar taxonomy in our analysis at the global level: section 3 focuses on growth and section 4 on preferences.

3. Global Growth and Real Interest Rates

One of the most frequently cited drivers of changes in real interest rates is changes in trend growth. Before analyzing how global trend growth might have changed, it is worth dwelling on how changes in growth affect real rates. The Euler equation in the neoclassical model pins down the real rate by time preferences, the pace of technological progress, and, in some formulations, population growth (equation (1)).

$$r^* = q/\sigma + \theta + (\alpha \cdot n), \tag{1}$$

where $r*$ is the real interest rate consistent with inflation at target and zero output gap in the long run; σ is household's intertemporal elasticity of substitution in consumption (preference for smoothed consumption); q is the rate of labor-augmenting technological change; θ is household's rate of time preference (patience); n is the rate of population growth; and α is the coefficient on the rate of population growth.[8]

The preference parameters are particularly noteworthy, as they affect the link between growth and real rates. Estimates of these parameters are difficult to obtain, but one meta-study by Havranek et al. (2015) suggests that the global average for σ (household preferences for smoothed consumption) could be around 0.5. This implies

[8]The infinite-horizon representative-agent Ramsey model does not include population growth in the steady-state real rate formulation. But there may be good reasons to include it (e.g., see Baker, Delong, and Krugman 2005).

that the mapping between productivity growth and real rates may not be one-for-one, but potentially one-for-two, i.e., a 1 pp fall in productivity growth could lead to a 2 pp fall in real rates.

The link between population growth (n) and real interest rates is even less certain. The standard neoclassical model does not include population growth as a driver of real interest rates at all (so $\alpha = 0$). Yet excluding population growth entirely from the analysis seems an omission. Alvin Hansen's (1934, 1938) original secular stagnation hypothesis emphasized the role of population growth in driving down the rate of return on capital. Baker, Delong, and Krugman (2005) have since argued along similar grounds, noting that in some models (such as the Solow model) population growth plays a role in determining real interest rates. If labor and capital are complements, then slower population growth should reduce the marginal product of capital, as firms have fewer workers to get the best out of their machines. As a result, slower population growth should mean the rate of return on capital falls, pushing down real interest rates. Given the lack of empirical estimates for α, we assume there is some role for population growth in driving real interest rates but a one-to-one mapping is likely to be an upper bound (i.e., $\alpha \leq 1$).

It is difficult to account for much (if any) of the pre-crisis fall in real rates by just appealing to past changes in growth, because global growth was fairly steady in the pre-crisis decades—averaging 3 to 4 percent per year. However, the crisis itself may have triggered a broader reassessment of trend growth expectations. And greater pessimism about future growth could be playing an important role in driving the decline in real rates we have seen most recently.

Broadly speaking, there are three factors that might lead trend growth to weaken over the future: (i) a reduction in labor supply growth; (ii) a slower rate of catchup in emerging markets; and (iii) weaker growth at the technological frontier. So how pessimistic should we be about each?

3.1 Labor Supply

Growth in global labor supply peaked at just over 2 percent in the 1980s as the demographic dividend from the post-war baby boom (and falling mortality rates in emerging market economies) fed through to the labor market. Since then, the pace of global

Figure 2. Global Population Growth and Labor Supply Growth

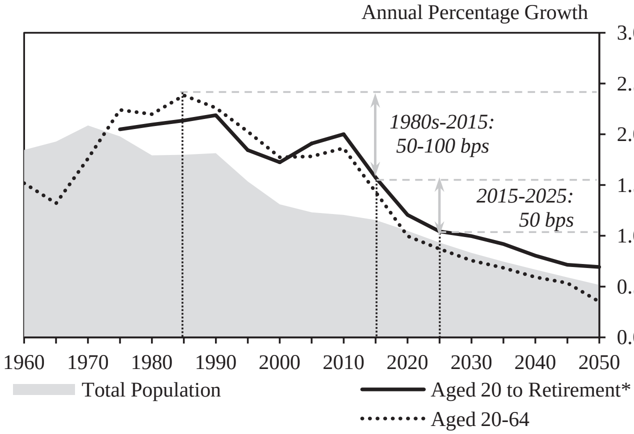

Notes: The solid black line shows growth in the global working-age population, aged twenty to retirement, using historical and forecast figures from the United Nation's Population Projections. Retirement ages are calculated using OECD data on average effective retirement ages. Over the future, global retirement ages are assumed to grow by one and a half years every decade to keep pace with expected rises in longevity.

labor force growth has slowed by a third, reducing its contribution to global GDP growth by around 1 pp (figure 2).

The age structure of the population means further falls are likely: global population growth slowed sharply in the mid-1990s, and that effect is now feeding through to global labor supply (figure 2). Expectations of slower labor supply growth could reduce global growth by around 0.5 pp over the next decade. As noted earlier, the mapping from labor force growth to real rates is highly uncertain (α in equation (1) is unknown). If we assume a mapping of one-to-one as an upper bound between labor force growth and real rates (i.e., $\alpha \leq 1$), this would suggest that expectations of slower labor force growth could account for up to 50 bps of the fall in real rates we have seen.

Slower labor supply growth could also have wider effects on economic growth and real rates via productivity spillovers—boosting productivity by alleviating resource constraints or reducing productivity by reducing the returns to innovation (see Kuznets's seminal

1960 paper for a discussion). However, given that the effects can go both ways, we do not explicitly account for these additional spillovers here.

3.2 Catchup Growth

Another driver of global growth is productivity catchup. As countries accumulate more capital and improve efficiency by adopting the latest technologies from overseas, productivity per worker rises. In the early 2000s, rapid catchup growth among Asian economies boosted global growth significantly, offsetting the impact of slowing labor force growth.

Given that the pace of catchup has accelerated in recent years, it seems implausible that slower convergence has been a key driver of the steady decline in real interest rates we have seen over the past few decades. That said, the robust rate of catchup growth seen in the early 2000s is relatively unusual when compared with a broader sweep of history (figure 3). Between 1980 and 2010, GDP per capita growth in the United States (widely used as a proxy for the technological frontier) was actually faster than the average across the rest of the world in fifteen out of thirty years—so the rest of the world spent just as long falling further behind the frontier as catching up.

The difficult question is what will happen to the pace of catchup going forward. In order for the recent positive trend to continue, many emerging markets will need to overcome the middle-income trap and continue to avoid geopolitical and financial crises. There are some grounds for optimism on that front. The unusually rapid growth of many emerging markets earlier this century shows that it is possible for sustained periods of catchup to take place. In addition, the rising importance of digital technologies in driving innovation, combined with the spread of the Internet and other communication technologies (e.g., distance learning), means it is now easier for ideas and skills to be shared across borders more quickly. On the other hand, the mixed performance of the 1980s and 1990s, combined with ongoing concerns about the stability of China's financial system and the rise of emerging market indebtedness, suggests that a more gloomy outlook is possible (at least in the near term).

We take a neutral view and assume that the contribution of catchup growth to global growth remains stable at its average rate

Figure 3. GDP per Capita by Region Relative to the United States

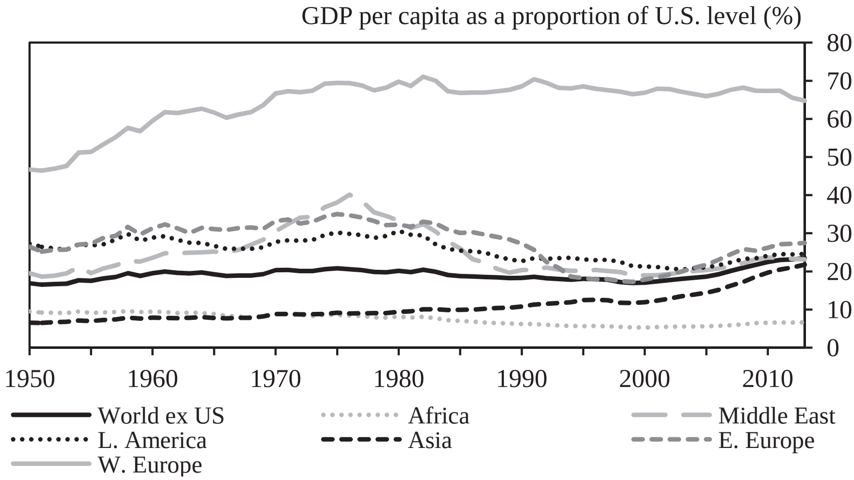

Note: Figures are calculated from the Conference Board's Total Economy Database (2013).

of the past twenty years—so not as fast as the early 2000s, but not as slow as the 1990s either. This equates to catchup growth continuing to add 1 pp to global growth per year and means we do not think the decline in global real interest rates is driven by slowing expectations of the pace of convergence.

3.3 Growth at the Technological Frontier

The other driver of global trend growth is the pace of growth at the technological frontier (proxied by productivity growth in the United States). Gordon (2014a, 2014b) has championed the view that several structural headwinds will hold back U.S. growth in the future, including further falls in the pace of educational attainment, rising inequality, and fiscal drag. Gordon suggests that these factors could drag down trend growth at the frontier by up to 1 pp—either by directly affecting the supply side or via weaker demand leading to hysteresis. Having interrogated Gordon's analysis, the argument on the educational plateau seems justified—the number of years of schooling per worker cannot go on rising indefinitely and has already started to slow. Moreover, Gordon's figures are backed up by others, e.g., Jorgenson and Vu (2010). But Gordon's estimates of the impact from inequality and fiscal policy seem high given uncertainties over fiscal multipliers and the overlap between the inequality and education arguments (rising inequality makes it more difficult for poorer households to afford college). In our view it seems more likely that growth at the frontier could be 0.5 pp weaker in the decade ahead due to these headwinds, i.e., half Gordon's estimate—Rachel and Smith (2015) contains the detailed analysis looking at the magnitude of each headwind that supports this claim. Given that the multiplier between productivity growth and real rates probably lies between 1:1 and 1:2 (i.e., $0.5 \leq \sigma \leq 1$ in equation (1)), expectations of slower frontier growth of around 0.5 pp could be pulling down on real rates by 50–100 bps.

The other major uncertainty is over the pace of innovation. Gordon (2014b) argues that the recent weakness in U.S. productivity growth is a longer-lived phenomenon stretching back to the 1970s and hence will continue going forward. Others, such as Brynjolfsson and McAfee (2014), see the recent slowdown as a blip—growing pains as a result of disruptive new digital technologies that will soon give way to rapid productivity gains. Some, such as Kurzweil (2005) even argue we are about to enter a phase of unprecedented growth, fueled by artificial intelligence and robotics. Our reading of the above arguments is that Gordon's characterization of recent history and the near future is the most compelling. Measured U.S. productivity growth has been weak since the 1970s; it was lifted temporarily by the information and communication technology (ICT) boom, but has since fallen back to its sluggish underlying rate. In the absence of clear advances in technology, it seems reasonable to assume this trend will continue in the near term—particularly given the recent weakness in productivity globally. Consequently, we assume that the pace of innovation will continue at its recent sluggish rate—consistent with the experience of the past thirty years, save for the ICT boom. However, a word of caution is warranted here. The further we peer into the future, the more likely a positive technology shock is to occur, so there are substantial upside risks to our forecast—Brynjolfsson and McAfee (2014) or even Kurzweil (2005) could eventually be proved right.

Figure 4 and the accompanying table bring together the above insights into one place—showing the contributions to each of the three factors to global growth. Looking over the past, it is reasonably clear that growth has not changed materially since the 1980s—averaging 3 to 4 percent per year. In fact, growth in the years before the crisis was actually a little higher than in the 1980s. Consequently, it is difficult to account for much (if any) of the pre-crisis fall in real rates by appealing to significant changes in the pace of global trend growth. Yet while growth may not have fallen much over the past, there are reasons to think global trend growth will slow in the future. Although there is a great deal of uncertainty, if we add up all the factors analyzed above, we think we can come up with a reasonable case for why global growth could slow by up to 1 pp over the next decade or so—largely relating to slower labor supply growth and structural headwinds at the technological frontier. Some of these factors could have been predicted before the financial crisis, but arguably the crisis was the trigger that caused financial markets to focus attention on the issue. Indeed, the 1 pp decline in global growth we predict is similar in scale to the downward revisions to medium-term growth forecasts that the International Monetary Fund (IMF) and private-sector forecasters have made since 2008. Depending on the mapping, these weaker growth prospects could account for around 100 bps of the fall in real rates we have seen post-crisis.

4. Desired Saving and Investment and Real Interest Rates

The previous section explored the link between global trend growth and the global neutral rate of interest. But as we have seen, changes in global growth can only explain part of the secular decline in global real rates over the past thirty years—mainly in the post-crisis period. Other factors must also be responsible for driving the long-term decline in the global neutral rate.[9] Since the real interest rate is the price of future consumption expressed in terms of consumption today, shifts in time preferences that describe how households

[9]Hamilton et al. (2015) arrive at a similar conclusion.

Figure 4. Global Growth Accounting

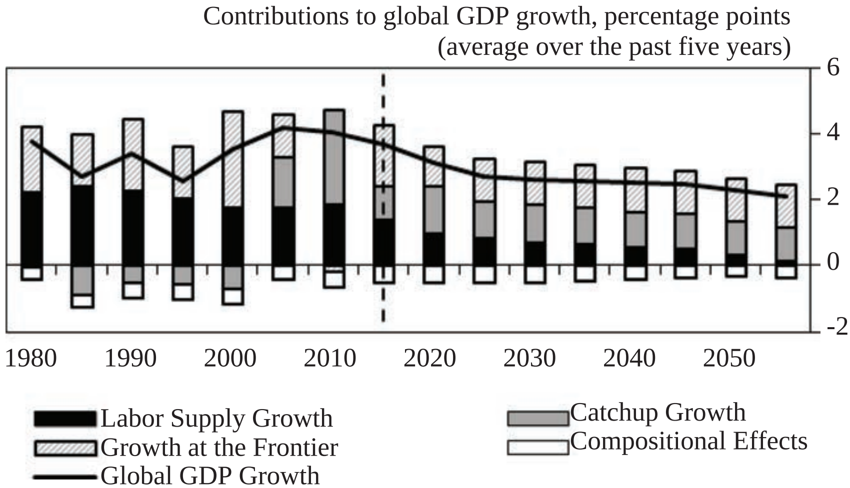

	1980 to 2015	2015 to 2030
Change in Global Growth	0.0 pp	–1.0 pp
Labor Supply Growth	–0.8	–0.5
Catchup Growth:	+1.0	0.0
Growth at the Frontier:	–0.2	–0.5
Educational Plateau	–0.2	(0 to –0.2)
Inequality	0.0	(0 to –0.6)
Fiscal	+0.2	(0 to –0.2)
Technological Progress	–0.2	(–)

Notes: Global growth is expressed in constant PPP-weighted 1990 dollars. Figures are expressed as averages over five-year intervals so as to smooth out cyclical variations and to provide a simple metric for the underlying trend. The white "compositional effect" bars in the figure show the impact on average global per capita incomes of having high population growth in low-income countries. Figures are constructed using data from the U.S. Conference Board's Total Economy Database, IMF, UN, and the authors' own calculations. Figures in parentheses show the range of uncertainty of the impact of each headwind on productivity growth at the frontier. Our central assumption is that the effects of the education plateau are towards the upper end of the range, while the fiscal and inequality effects are towards the lower end. Global growth is expected to slow by up to 1 pp weaker over the next decade due to slower labor supply growth (–0.5 pp) and slowing productivity growth at the frontier (–0.5 pp). The multiplier from changes in growth to real rates is likely to range from 1:1 to 1:2 for productivity growth and 1:0 to 1:1 for labor supply growth. This implies that slowing growth expectations could have reduced real rates by between 50 and 150 bps. We quote the midpoint of this range in our calculations above.

spread consumption over their life cycle will also move real rates around. The simple growth model we introduced in the previous section is too parsimonious to permit the analysis of the kind of preference shifts we are concerned with here: the secular trends considered below would show up, in a reduced-form way, as changes in the value of the parameters of the Euler equation, and would not be very informative about the workings of various mechanisms or their magnitude. So instead of a formal model, in this section we utilize a saving-investment (S-I) framework to shed light on these phenomena.

The basic idea behind this framework is that, given growth expectations, the neutral rate will depend on agents' preferences for desired saving and desired investment. Intuitively, desired saving will tend to rise as real rates increase (the saving schedule should slope upwards), because higher rates generate higher returns on saving and yield higher future consumption. By contrast, desired investment will tend to fall as real rates rise (the investment schedule should slope downwards) because the real rate is a key component of the user cost of capital, so as real rates rise it becomes more costly to invest. The focus of our analysis is on changes in *desired*, rather than *actual* saving and investment. For the world as a whole—as for any closed economy—actual saving and investment will always be equal by identity. But the sensitivity of desired saving and investment (the slopes of the curves) and the forces that shift them (preference shifts) will be key in determining the actual (equilibrium) level of saving and investment and the observed real interest rate.

An important source of uncertainty when using this framework for quantitative analysis is that we do not know the sensitivities of desired saving and investment to real rates: the slopes of the saving and investment schedules are unobservable, so we need to rely on empirical estimates. The key difficulty when estimating these slopes is endogeneity: interest rates and S-I ratios may be driven by common factors. For example, a more optimistic demand outlook would raise investment and interest rates simultaneously. This is why studies that estimate elasticities using time-series correlations can produce a wide range of estimates. To make our exercise robust, we take an average of available estimates from the literature and then later conduct sensitivity analysis to show the impact of varying these assumptions.

Table 1. Estimates of the Elasticity of Saving with Respect to the Real Rate

Author of Study	Elasticity
Blinder (1975)	0.03
Boskin (1978)	0.4
Carlino (1982)	0
Carlino and DeFina (1983)	0
Gylfason (1981)	0.3
Heien (1972)	1.8
Howrey and Hymans (1978)	0
Summers (1982)	1.3
Taylor (1971)	0.8
Wright (1967)	0.2
Average	0.5

Source: DeFina (1984).

For the elasticity of saving, this suggests an elasticity of 0.5 (table 1), although the range of estimates is admittedly very wide. For the elasticity of investment, we rely on more recent studies that aim to overcome the endogeneity problem by using structural models or by employing cross-sectional data (such as Ellis and Price 2003, Gilchrist and Zakrajsek 2007, and Guiso et al. 2002). These tend to find that long-run elasticities are between –0.5 and –1. We assume an elasticity of –0.7; this makes investment more sensitive to interest rates than saving.

Together, these assumptions form the basis of the slopes of the curves shown in figure 5. We think the slopes have been calibrated based on a fairly neutral reading of the range of estimates in the literature, but we recognize the wide bands of uncertainty. If we are wrong about one of the slopes of the curves—say the investment schedule is shallower—then it becomes more likely that shifts in the investment curve (rather than the saving curve) have been responsible for more of the fall in real rates we have seen. The slopes of the curves thus matter in terms of the relative weight one puts on different explanations for the fall in real rates, but should not necessarily affect our ability to account for the scale of the fall overall.

Figure 5. The Saving-Investment Framework

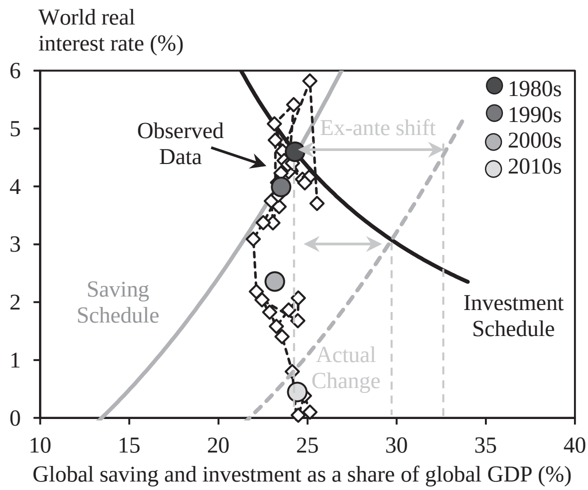

Notes: Constructed using data from the IMF, King and Low (2014), and authors' calculations. Global saving and investment rates are reported by the IMF. The world real rate is taken from King and Low (2014). Specific calibrations of neoclassical or overlapping-generations models can deliver very different results.

A more detailed discussion of the sensitivity of the results to these assumptions is provided at the end of this section.

We choose the average level of real rates during the 1980s as the reference point of our analysis. This choice is driven by the observation that the initial decline in real rates from the peak of around 6 percent in the early 1980s was likely driven by the disinflation policies during Volcker's Chairmanship of the Federal Reserve Board (IMF 2014b). The average of the 1980s also equates to the level of real rates in the first half of the 1990s, when such monetary policy regime shifts had largely been completed.

We focus on ex ante shifts in the schedules, i.e., shifts that are independent of the moves in the real rate itself. We can then determine the impact on the real rate by comparing the new and old equilibria. Precisely because the real rate adjusts, the ex post change in actual saving or investment will tend to be smaller than the ex ante, or desired, change. Figure 5 illustrates this with an example: a preference shift increases desired saving for any given interest rate.

This shifts the saving schedule to the right. But as the desired level of investment is unchanged (for a given interest rate), this shift would push down on the interest rate until desired investment is equal to desired saving. As a result, the actual increase in saving is smaller than the shift in desired saving.

One striking feature of the data is that despite the 450 bps fall in global real rates, global saving and investment have remained fairly stable as a share of global GDP over the past thirty years (diamonds in figure 5). This vertical pattern could suggest that either saving or investment is insensitive to changes in real rates (one of the curves is vertical). While mindful of this possibility, we assume that the slopes of the curves match empirical estimates in the literature, which implies that *both* curves must have shifted. Various factors have been put forward to explain such shifts. Our approach is to run through them and try to quantify the size of each effect on real rates. We begin by focusing on trends that have affected the saving schedule: changes in the demographic structure of the global population, rising inequality, and a preference shift by emerging market governments towards higher saving (the emerging market saving glut). We then analyze trends that have mostly affected desired investment: the fall in the relative price of capital, shifts in public investment, and the changes in the spread between broad rate of return to capital and the risk-free rate.

We view the estimates below as largely independent of the growth effects we have identified in the last section, and hence claim that our results should not suffer from double-counting. Although we do not have a comprehensive general equilibrium model to prove this point, we rely on the economic logic that is commonly accepted in the macroeconomic literature: namely that in macroeconomic equilibrium, the response of the real interest rate will cushion the impact on economic growth. Up to the most recent crisis, the real interest rate was falling, but it remained above zero. Thus, for most of the period under study, the equilibrium rate had sufficient room to respond, and precisely because of that response the secular forces we discuss should not have had much impact on the growth rate up until recently. Take inequality as an example: given a fixed real rate of interest, rising inequality could have lowered aggregate demand and pushed down on the growth rate. But as long as there was room for the interest rate to fall, this direct effect would have been offset

Figure 6. Cross-Country Saving Rates and Dependency Ratios

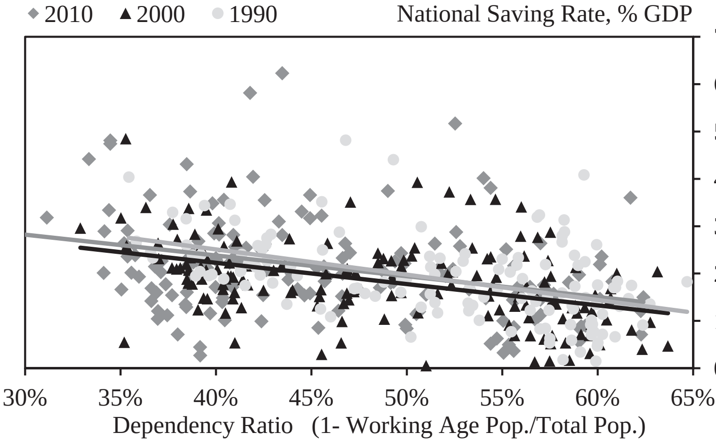

Sources: IMF's *World Economic Outlook* and UN Population Statistics.

Notes: The scatter plot shows cross-country values of dependency ratios and saving rates for three different decades. Working-age population is defined simply as those aged twenty to sixty-five.

in equilibrium. Confirming this intuition, the moving average of the global growth rate appears to have been relatively stable over much of the time period we analyze, as highlighted in section 3.

4.1 Demographics

The key mechanism through which demographics plays a role in our analysis is the life cycle of saving. Consumption is fairly stable over the life cycle, but income is hump shaped, so people of working age are those who tend to save the most. Consequently, the greater the proportion of the population that is of working age, the higher the desired level of saving in aggregate is likely to be.

This simple intuition is discernible in cross-country data on national saving rates (figure 6). There is a significant negative relationship between the dependency ratio (defined as the proportion of the population not of working age) and national saving rates: every 1 pp fall in the dependency ratio translates to around a 0.5 pp rise in national saving rates. This relationship is stable through time,

Figure 7. Shift of the Desired Saving Schedule Due to Changes in the Composition of the Global Population

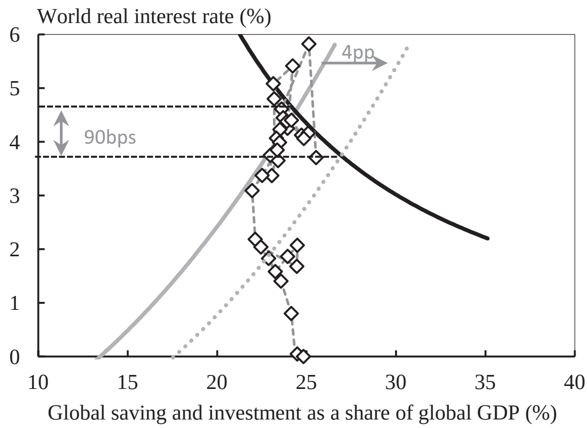

Sources: IMF and authors' calculations.

suggesting that it is robust and can be used to calibrate the ex ante shift of desired saving caused by demographic changes at the global level.

Over the past thirty years the proportion of dependents has fallen from around 50 percent of the global population to 42 percent. The main driver of this decline has been a fall in the proportion of young dependents—reflecting the slowdown in demographic growth discussed earlier. This effect has more than offset the gradual rise in the proportion of old-age dependents linked to aging societies. Using the estimated cross-country relationship depicted in figure 6, the 8 pp fall in the global dependency ratio translates to a 4 pp rise in desired saving, for a given real interest rate.

The 4 pp rise in desired saving can be illustrated by a rightward shift in the saving schedule (figure 7). The effect on the global real rate can then be easily read off the y-axis of the chart by comparing the two intersection points. This suggests that the effect of the fall in the dependency ratio has been to lower real rates by around 90 bps.[10]

[10] Aside from changes in the global dependency ratio, changes in average age of the working-age population may also shift desired saving. This is because over

Looking ahead, the dependency ratio is likely to stop falling or even rise as a growing share of the world's population enters retirement age. At face value this suggests that the effects of demographics on desired saving and hence real rates could reverse. But the extent of this reversal is uncertain. Two factors that may limit the extent of the reversal are increases in retirement ages and longevity. Effective retirement ages across countries have been increasing gently over the past fifteen years. This trend largely reflects changes in old-age participation rates rather than official retirement ages, which have been fairly static. If this trend continues, the average retirement age across the OECD could reach sixty-seven by 2030—enough to halt any uptick in the dependency ratio. In addition, an increase in longevity over and above this increase in retirement age may mean that people of working age choose to save a larger share of their income to fund a longer retirement—pushing up on saving rates (Carvalho, Ferrero, and Nechio 2016). These two factors suggest that there could be little to no reversal of the impact of demographic forces on global real rates in the years ahead.

On the other hand, the shifting composition of dependents from young to old could mean that the reversal is larger than suggested by simple dependency metrics. Old-age dependents can have much lower net saving rates than young dependents due to the distinct pattern of consumption over the life cycle. In advanced economies (such as the United States), consumption tends to drift up in retirement, particularly in the last few years of life. This is driven largely by greater consumption of health care (hospitalization, emergency procedures, etc.).[11] Thus in advanced countries, a rising share of retirees will have a disproportionately large *negative* effect on desired saving—implying a faster turnaround in the impact of demographic trends on desired saving globally compared with the simple dependency ratio metrics. Indeed, studies that look at the "support ratio"—which can be thought of as a measure of the dependency ratio that corrects for the differential spending patterns highlighted

their working lives, individuals' saving rates tend to rise up to a certain age and then decline beyond that. Our analysis suggests that this effect is second order compared with the 4 pp shift due to the dependency ratio discussed above. See Rachel and Smith (2015) for a fuller discussion.

[11]This stands in contrast to a more flat pattern in emerging countries, such as China—see Lee and Mason (2011) for more details.

above—find that the impact of demographics on desired saving looks set to reverse (Lee and Mason 2011).

On balance, we think that the evidence based on support ratios provides the more reliable steer, and so we expect the demographic effect on rates to gradually reverse. The current projections of dependency ratios alone would suggest that demographics will be neutral on global real rates going forward, but given that older people in advanced economies are big spenders, desired saving is likely to decrease. Quantitatively, the current forecasts for the support ratios in advanced and emerging economies suggest that the reversal of the demographic effect over the next twenty years is likely to be about half that of the downward drag on real rates over the past twenty years.

4.2 Inequality

Changes in the distribution of income can affect desired saving because the rich and poor tend to save different proportions of their income. To the extent that the rich save more, rising inequality will result in lower consumption, higher desired saving, and hence a lower equilibrium real rate.[12] Empirical evidence supports the notion that the rich do save more, although the range of estimates available is relatively limited and primarily covers the United States.[13] The seminal contribution is by Dynan, Skinner, and Zeldes (2004), who show that average saving rates and marginal propensities to save tend to rise with the level of income. This is confirmed by more recent evidence: for example, Cynamon and Fazzari (2016) show that the richest 5 percent save much more than the rest (with saving rates around three times as high), and Saez and Zucman (2014) give a long-run perspective on the high saving rates of the wealthy.

Two global trends have taken hold since the 1980s: inequality *between* countries has fallen because developing countries, particularly those in Asia, have been catching up with their Western peers; but at the same time income inequality *within* countries has been rising (Piketty 2014). The relationship between national saving rates

[12]This general mechanism has been incorporated in recent models, e.g., in Eggertsson and Mehrotra (2014) and Kumhof, Ranciere, and Winant (2015).

[13]Crossley and O'Dea (2010) provide a UK perspective.

and a country's level of development is not clear cut—for example, many Asian economies have very high saving rates despite having relatively low incomes. As such, it is less clear what impact the reduction in inequality between countries has had on global saving rates. By contrast, the relationship between individual saving rates and individual income levels within countries is better established. We therefore focus on calibrating the impact of this latter effect on desired saving rates globally.

To isolate the effect of rising inequality within countries on desired saving, we perform a thought experiment using U.S. data. First, we take the saving rates from Dynan, Skinner, and Zeldes (2004) for different income quintiles. Then we combine those saving rates with data from the U.S. Census Bureau showing how income shares across the population have changed since the 1980s. Over this period, the richest fifth of the population—who are also the keenest savers—have seen their share of national income rise by around 7 pp. On average, this group saves an extra third of their income compared with the rest of the population, so this shift in the income distribution translates to a rise in desired saving of around 3 pp. Changes in the income distribution among the four lower quintiles of the population reduce this figure to a net rise of around 2 pp in aggregate. This acts to drive the real interest rate downward by 45 bps.[14]

Our assessment of the effect of inequality relies primarily on U.S. data, which could overestimate the size of the effect at the global level, since the rise in inequality in the United States has probably been a little larger than average. But there are also other reasons why inequality may have had a bigger effect on real rates than postulated above.

First, over the past thirty years, the share of global income earned from owning capital has risen as the labor share has fallen. Mechanically, this should push up on desired saving if the propensity to save out of capital income is higher than from labor income. A lack of empirical data on saving rates from these different sources of income prevents detailed analysis of this channel, but we note it as an upside risk.[15]

[14] We provide further cross-checks on this figure in Rachel and Smith (2015).

[15] See Rachel and Smith (2015) for further discussion.

Second, inequality between countries has been falling as real incomes in the bottom two-thirds of the global income distribution have risen rapidly. The impact of this income shift on global saving depends on relative saving rates between lower-income countries that are catching up and advanced economies. National saving rates in advanced economies and emerging economies were virtually identical up until the year 2000, so the effect on global desired saving of emerging markets catching up would have been negligible until then. But since the turn of the century, saving rates in emerging markets have actually increased *above* those in advanced economies. This means that faster income growth in emerging market economies may have raised global desired saving.[16] We are cautious of including this effect separately, as we think it is at least partly driven by cyclical factors and partly related to the emerging market saving glut story, which we discuss next. To avoid double-counting, we exclude it from our analysis.[17] Instead, we focus on the future of inequality within countries and how that affects global saving.

The future of within-country inequality will ultimately depend on policy. Piketty and Saez (2014) point out that "inequality does not follow a deterministic process. In a sense, both Marx and Kuznets were wrong. There are powerful forces pushing alternately in the direction of rising or shrinking inequality. Which one dominates depends on the institutions and policies that societies choose to adopt." So any forecast of inequality will necessarily rely on heroic assumptions about inherently unpredictable political processes. That said, it may still be useful to consider the main *economic* forces that determine income inequality in the long run, taking policies and political processes and institution as given.

The most widely used model for analyzing *labor income inequality* is based on the idea of a race between education and technology (Goldin and Katz 2007). The basic intuition is that more education leads to a rise in the supply of skilled labor, while technological change leads to a rise in the demand for skilled labor. If there is a relative shortage of high-skilled labor, because technological progress

[16]Buiter, Rahbari, and Seydl (2015) makes a similar point.

[17]Taken literally, the rise in emerging market economies' share in world GDP would mechanically push up on desired saving by around 2 percent of world GDP. The effect is strongest since 2007. The pickup is largely driven by China.

races ahead of educational attainment, then inequality is likely to increase, as those with the sought-after skills will tend to see their earnings rise relative to the rest of the population. In this context, the recent rise in the cost of education, if not tackled by policy, could potentially represent a rise in inequality of opportunity that may limit educational attainment and hence increase income inequality, probably with a long lag. This mechanism could be further strengthened if the latest technological advances not only increase the demand for skilled workers but also replace low-skilled jobs—Frey and Osborne (2013) predict that 47 percent of U.S. employment may be subject to computerization over the next twenty years.

Other factors may also play a role. For example, further globalization could make "winner-takes-all markets" more common, raising the share of income accruing to the global "superstars." Piketty (2014) also suggests that if the growth rate of labor income declines as global growth falls back (g), but the rate of return on capital (r) is maintained at its historic rate, then inequality is likely to rise further. On balance, absent a major policy shift, we judge that labor income inequality is more likely to continue rising than to fall back in the years ahead, but the future path of inequality is very uncertain. Hence our treatment is cautious, assuming only a very gentle increase in inequality going forward.

4.3 The Emerging Market Saving Glut

Following the Asian crisis in 1998, many emerging markets significantly increased their foreign exchange reserves as a precautionary measure against the future risk of destabilizing capital outflows. In tandem, the era of high oil prices prompted an increase in saving among oil producers. Bernanke (2005) suggests that these forces represented a preference shift by governments (in Asia) and a shift in circumstances (for oil exporters) that were largely exogenous to the global system. These preference shifts resulted in an increase in desired saving in those countries. To the extent that this increase was not matched by a rise in desired investment, it led to a net increase in global saving.

On average, the current account surplus of Asian economies and oil exporters—indicative of the net amount of financial capital that those countries send abroad—has been 1 percent of world GDP since

Figure 8. Global Imbalances

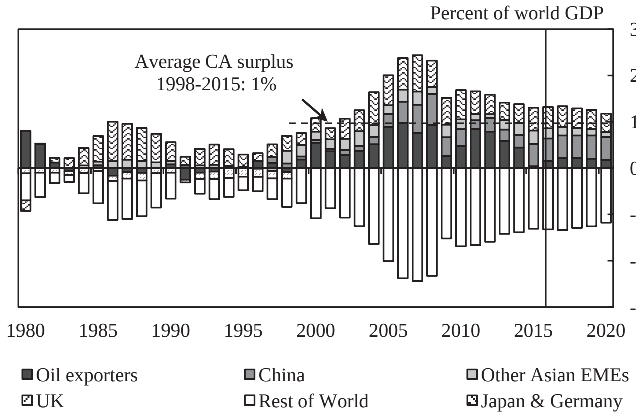

Sources: IMF and authors' calculations.

the late 1990s, around 1 pp higher than the roughly balanced current account pre-1998 (figure 8). Using the increase in emerging markets' current account surplus as a guide suggests that the desired saving schedule has shifted to the right by 1 pp as a result of the emerging market saving glut, which lowers the global real rate by round 25 bps. This is only around half of the effect of inequality, and a quarter of the effect of demographics.

Bernanke (2015b) discusses the future of the emerging market saving glut and concludes that the outlook is mixed. On the one hand, three factors suggest that these imbalances may have run their course: (i) some of the emerging market economies, particularly China, are rebalancing their economies away from exports toward domestic demand; (ii) the buffer stock of FX reserves that emerging markets hold is already sufficiently large, and the buildup of foreign currency reserves is slowing and in some cases now falling; and (iii) oil prices have fallen, so we might expect the excess savings from oil producers to decline further from pre-crisis peaks. On the other hand, there are also some new potential sources of the "saving glut." For example, Bernanke notes that Germany has the highest current account surplus in the world, and there is a concern that this will persist, exerting further downward pressure on global real rates. But to us it is unclear whether Germany's surplus will act as an additional force or has already been captured by other trends discussed in this paper, notably demographics. Overall, we think that the IMF forecast for global imbalances—as shown in figure 8—is a reasonable baseline forecast, which suggests a very gradual unwinding of the emerging market saving glut going forward.[18]

Taken together, shifts in desired saving linked to the three trends above can account for around 160 bps of the fall in global real rates since the 1980s. If this had been the whole story, we would have expected to see a steady rise in actual saving rates globally. But global saving and investment ratios have been remarkably stable over the past thirty years—as noted earlier. This suggests that the desired investment schedule has also shifted. Here we focus on three trends that could potentially explain such a shift: (i) the secular decline in the relative price of capital goods; (ii) a preference shift away from public investment projects; and (iii) an increase in the spread between the risk-free rate and the return on capital.

4.4 The Falling Relative Price of Capital Goods

Perhaps one of the most pervasive trends that may have affected desired investment expenditure is the 30 percent decline in the relative price of capital goods since the 1980s.[19] Cheaper capital means that a given investment project costs less to pursue, so investment volumes can be maintained by committing a smaller share of nominal GDP. But cheaper capital also incentivizes additional investment projects, given the lower cost. The overall impact on capital expenditure is the sum of the two effects—its sign depends on the elasticity of substitution between capital and labor.

If capital and labor are easily substitutable, a fall in the relative price of capital goods will induce a lot of additional investment

[18]Others present compelling alternative forecasts. For example, Speller, Thwaites, and Wright (2011) perform long-run simulations of gross capital flows out to 2050 and predict persistently high current account surpluses in emerging markets, with the Chinese current account shrinking but India's rising.

[19]IMF (2014b) has examined changes in the relative price of investment in the advanced economies since 1980. The Fund documents a downward trend in the relative price of investment that then levels off in the early twenty-first century. In explaining this movement, the Fund points to the work of Gordon (1990), who emphasizes the role of research and development that is embodied in cheaper, more efficient investment goods.

projects, potentially by enough to counter the effect of falling prices and hence maintain investment as a share of nominal GDP. But most empirical work points to the elasticity being smaller than one.[20] The IMF (2014b), for example, asserts that any increase in the volume of investment caused by a decline in the price of capital goods has been insufficient to offset the negative impact on real interest rates. Thwaites (2015) surveys the literature and arrives at a similar conclusion. An elasticity smaller than one means that a fall in the relative price of capital goods will tend to be associated with a shift of the investment schedule to the left (desired investment expenditure is lower for a given interest rate). To calibrate the size of this shift, we rely on the model developed in Thwaites (2015). A 30 percent decline in the relative price of investment lowers the steady-state nominal investment-to-GDP ratio by around 1 pp in that model. The fall in the relative price of capital goods also has an additional effect, which is to pivot the investment schedule (so that it becomes steeper), as any given amount of real investment now requires less of today's output to be sacrificed. In other words, the opportunity and financial cost of investment become a less important factor in making investment decisions. Desired investment becomes less sensitive to interest rates by roughly the same amount as the fall in the relative price of capital goods.[21]

The 1 pp shift in the investment schedule, together with a 30 percent drop in the elasticity of investment with respect to the interest rate, delivers around a 50 bps fall in the real rate in the saving-investment diagram. This is similar to the peak-to-trough fall in the interest rate along the transition path to the new steady state in Thwaites (2015).[22]

The future of capital goods prices is still being hotly debated, and no clear consensus has yet emerged. Eichengreen (2015) argues that further falls are not guaranteed:

[20] Piketty (2014) is a notable exception.

[21] More formally, in the formulation of the user cost of capital, the weighted average cost of capital enters multiplicatively with the relative price of capital goods.

[22] The paper argues that the initially positive response of the interest rate matches what happened in the real world in the 1970s, when real rates were increasing. Since we are trying to explain the fall in real rates since the 1980s, it is appropriate to compare the decline in the real interest rate since the peak.

> Evidently, R&D is not embodied more easily and fully in investment goods than consumption goods in all times and places. The presumption behind "the Baumol effect"—that consumption goods, and in particular that portion provided by the service sector, are difficult to mechanize and therefore become relatively more expensive over time—may not hold in the future as it has in the recent past. Even if the post-1980 decline in the relative price of investment goods is part of the explanation for the concurrent decline in real interest rates, there is no ruling out that it may be reversed in the future.

Others have also noted that the relative price of capital goods has stabilized more recently, taking this as evidence that the ICT revolution has run its course.

On the other side of this debate, researchers at the Federal Reserve (Byrne, Oliner, and Sichel 2013) believe that the price of ICT equipment has been persistently mismeasured. In their view, statisticians struggle to capture the higher capability of the latest technologies such as quad-core processors. Byrne et al.'s quality-adjusted estimates suggest that microprocessor prices continue to fall at a rate of around 30 percent a year. Furthermore, Thwaites (2015) argues that the effect on real rates can build for a long time even after capital goods prices stop falling.

Overall, it seems reasonable to assume some further contribution from the decline in the relative price of capital goods to lower rates, albeit at a diminished magnitude compared with the past.

4.5 Lower Public Investment

Public investment has been on a declining trend as a share of global GDP since the 1980s (see IMF 2014a). This could be because political views have become more polarized, thus making it difficult to agree upon and implement large-scale public investment projects. Or it may be because voters have shifted their preferences away from large governments. Either way, the result of this shift has been to lower the global investment-to-GDP ratio by around 1 pp between 1980 and 2007. Since 2007, public investment in emerging economies—particularly China—has accelerated rapidly, unwinding the long-term decline seen in the preceding decades. However, we think that much of the post-2007 pickup is a cyclical response to weakening demand during the global financial crisis. We therefore expect this to reverse and the downward secular trend to eventually reassert itself. Consequently, we think that lower public investment has shifted the desired investment curve to the left by around 1 pp, lowering real rates by around 20 bps—a relatively small effect.

An alternative interpretation of the recent movements is that higher public investment in emerging market economies is currently pushing up on the global real rate (relative to pre-crisis), and if that unwinds, the global real rate will fall further. The difference between these two explanations comes down to whether the shift away from public investment has already affected global equilibrium rates or whether this is still to come. In either case, given that the size of the effect is relatively small, this channel is not a major driver of movements in the global real rate.

4.6 Spread between the Risk-Free Rate and the Rate of Return on Capital

So far our analysis has abstracted from the fact that the interest rate that matters for firms' investment decisions is the rate of return on capital, not the risk-free rate. Strictly speaking, when analyzing desired investment, the rate of return on capital rather than the risk-free rate should appear on the vertical axis of the S-I diagram. This distinction would not be important if the spread between the risk-free rate and the return on capital had been constant over time—the desired investment schedule shown in this paper would represent a simple vertical transformation of the "correct" schedule. However, there is some evidence that the spread has risen over time, which has implications for desired levels of investment. A rise in the spread shifts the desired investment schedule vertically down—because in order to keep desired investment unchanged, the risk-free rate must fall by exactly the same amount as the spread has increased, all else equal. However, in general equilibrium all else is not equal, and a lower risk-free rates induces people to save less—suggesting that the eventual decline in the risk-free rate may be a little smaller than the rise in the spread.

The IMF constructs a weighted measure of the spread across these measures for the world as a whole (figure 9). This shows that the rate of return on capital has fallen since the early 1990s, but

Figure 9. The Global Risk-Free Rate and Rate of Return on Capital

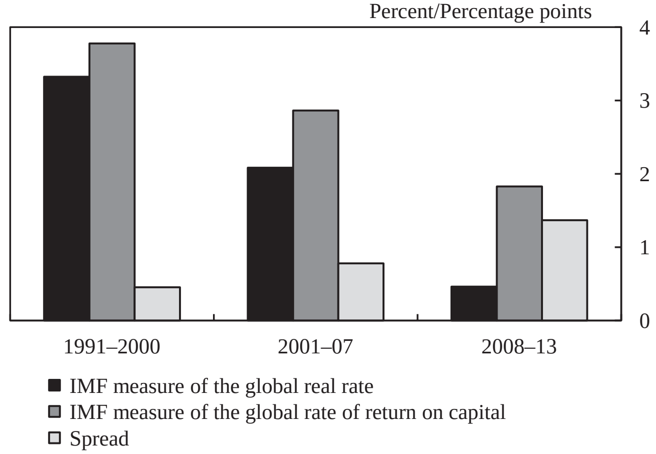

Sources: IMF (2014b) and authors' calculations.

not by as much as the risk-free rate—the spread has increased by around 100 bps. Market-by-market analysis supports this conclusion (Rachel and Smith 2015). In our framework, such a rise in the spread causes the real interest rate to drop by about 70 bps.

Together, the three investment trends highlighted above (the falling price of capital goods, lower public investment, and rising credit spreads) account for around 140 bps of the fall in real rates seen since the 1980s—a similar order of magnitude to the effect from the three saving trends highlighted earlier.

Figure 10 brings together all of these trends using the saving and investment framework outlined earlier. Our analysis suggests that the desired saving schedule has shifted out materially due to demographic forces (90 bps of the fall in real rates), higher inequality within countries (45 bps), and a preference shift towards higher saving by emerging market governments following the Asian crisis (25 bps). In addition, desired investment rates appear to have fallen as a result of the decline in the relative price of capital goods (accounting for 50 bps of the fall in real rates), a preference shift away from public investment projects (20 bps), and an increase in the

Figure 10. Quantifying Shifts in Desired Saving and Investment

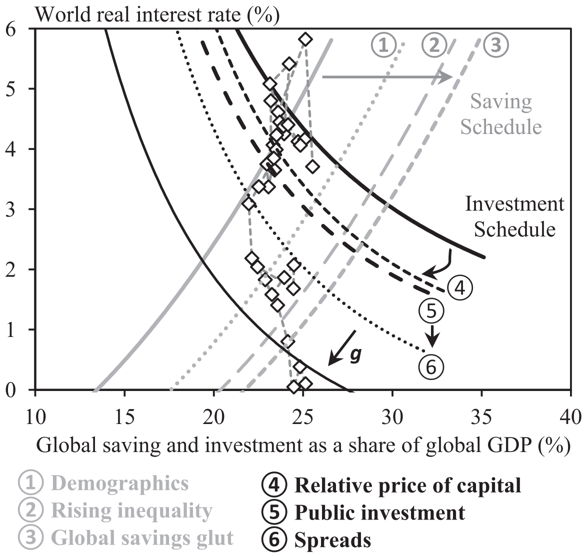

spread between the risk-free rate and the return on capital (70 bps). Together these effects can account for 300 bps of the fall in global real rates. We also include an illustrative shift of the desired investment schedule to account for weakening global growth prospects identified in the previous section (labeled "g" in the diagram).

This saving-investment framework provides a broad description of the relative sizes of the different forces at play. Taken at face value, shifts in preferences appear to explain around 300 bps of the decline in real rates since the 1980s, on top of the 100 bps explained by the deterioration in the outlook for trend growth seen more recently. In other words, we think we can account for most of the decline in global real rates using evidence independent of the decline itself.

Around 50 bps of the fall in real rates remains unaccounted for. This could reflect a number of factors, and each of these underscores the possible uncertainties that surround our point estimates. First, we might be missing certain secular trends from our analysis. For example, rising short-termism (Gutiérrez and Philippon 2016) or the decline in capital intensity of production (Summers 2013) could also be pushing down on real rates. Second, some of the trends we have quantified could be having bigger effects than we have estimated—for example, some studies that focus on an individual factor have found larger effects (e.g., for inequality and demographics). Third, the unexplained component could reflect global headwinds from the financial crisis, such as deleveraging or heightened risk aversion, which are temporarily pushing down on real rates. Fourth, the market measures of real rates we are using, which are derived from government bond yields, may be distorted. Post-crisis regulatory changes may have increased demand for safe government assets by financial institutions, while central bank quantitative easing (QE) has been temporarily boosting the demand for government bonds.

As mentioned earlier, significant uncertainty also surrounds the slopes of the saving and investment curves. The vertical pattern of the intersection between real interest rates and saving-investment ratios over the past thirty years may suggest that one or both of the schedules is steeper than we have assumed. While that is a possibility, we note that a near-vertical slope would imply that agents' behavior is invariant to changes in interest rates, and so would be at odds with empirical evidence suggesting that interest rates do affect behavior and the macroeconomy (Christiano, Eichenbaum, and Evans 1999; Ramey 2016). Moreover, it would likely lead to significant volatility in the observed interest rates, as small shifts would have a very large impact on the equilibrium real rate. This would be at odds with the relatively smooth downward trend in the long-term real rate over time. On the other hand, if both of the curves were significantly flatter than we assume, we would not have observed the large decline in the real interest rate, even in the presence of strong underlying secular trends: small adjustments in the interest rate would have realigned desired saving and investment. To check the robustness of our results, we vary the elasticities of the desired saving and investment curves in a number of ways. Table 2 reports four sensitivity tests compared with our benchmark calibration. The key point here is that while the slope of the curves affects the allocation of saving and investment trends in driving the change in real rates, when both sets of forces are viewed together they produce a broadly similar total impact on real rates across calibrations. Put differently, the conclusion that around 3 pp of the downward move in the real interest rate is due to shifts in desired saving and

Table 2. Sensitivity of the Results to Changing Elasticities of Saving and Investment Curves

Calibration	**Baseline**	**Saving More Sensitive to Interest Rates; Investment Less Sensitive**	**Saving Less Sensitive to Interest Rates; Investment More Responsive**	**Both Saving and Investment More Responsive (Less Steep)**	**Both Saving and Investment Less Responsive (Steeper)**
Elasticities	$\varepsilon_S = 0.5$; $\varepsilon_I = -0.7$	$\varepsilon_S = 1$; $\varepsilon_I = -0.35$	$\varepsilon_S = 0.25$; $\varepsilon_I = -1$	$\varepsilon_S = 1$; $\varepsilon_I = -1$	$\varepsilon_S = 0.25$; $\varepsilon_I = -0.35$
Shifts in Desired Saving	1.6 pp	1.7 pp	1.3 pp	1 pp	2.5 pp
Shifts in Desired Investment	1.4 pp	0.8 pp	1.6 pp	1 pp	1.7 pp
Total	3 pp	2.5 pp	2.9 pp	2 pp	4.2 pp

Note: ε_S and ε_I are estimates of the elasticities of saving and investment to the real interest rate, respectively.

Figure 11. Secular Drivers of Global Real Interest Rates

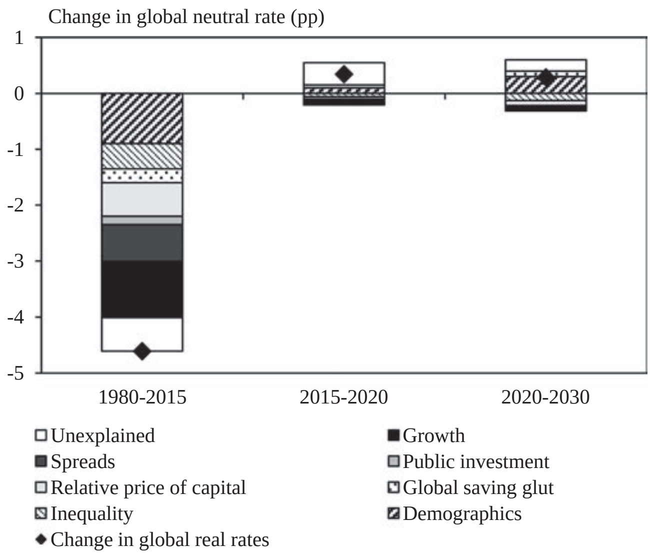

investment seems to be broadly robust to assumptions about the elasticities of the saving and investment schedules.

Clearly, the confidence interval around all of the above estimates is very wide. Nevertheless, quantifying the potential size of the impact from each of the different secular trends serves a useful purpose—it not only helps us explain movements in real rates over the past but also allows us to opine on how such trends (and hence real rates) are likely to evolve going forward. Figure 11 provides a summary of our findings about past movements in real rates, together with our best judgments about the direction of travel based on the discussions in this paper. The main message is that the trends we have analyzed are likely to persist at their current level: we do not predict a big further drag, or a rapid unwinding of any of these forces. Some are likely to drag a little further (global growth is set to decline further out, and we assume this will feed into slightly lower rates in anticipation; the relative price of capital is likely to continue to fall, albeit at a slower pace; and inequality may continue to rise), but this will be broadly offset by a rebound in other forces (particularly demographics). What happens to the unexplained component depends on what is driving it. In figure 11 we illustrate the effect of assuming that it is largely cyclical. Despite this, our predictions still imply that the global neutral rate will remain low, perhaps settling at around 1 percent in the medium to long term.

5. Conclusion

Since the 1980s, market measures of long-term risk-free real interest rates have declined by around 450 bps across both emerging and developed economies, with a discernible common trend suggesting that global factors are at work. The causes of the fall are likely numerous and diverse. In this paper we have attempted to quantify the impact that several secular forces could have had on real rates, by affecting global growth, desired saving, and desired investment. Although there is great uncertainty, our estimates suggest that these secular forces can account for about 400 bps of the 450 bps decline in the global long-term neutral rate seen since the 1980s. Moreover, most of these secular trends look likely to persist, suggesting that the global neutral rate may settle at around 1 percent over the medium to long run.

The policy implications of permanently low real interest rates are extensive. In the face of adverse shocks, central banks are likely to run up against the zero lower bound on nominal interest rates more often, requiring the use of unconventional policy instruments such as quantitative easing (QE). For large adverse shocks, fiscal policy may need to bear more of the burden of business-cycle management. Low rates may also fuel search-for-yield behavior, posing challenges for macroprudential and microprudential policymakers. More generally, the possibility of the global neutral rate remaining at persistently low levels should motivate a real debate across the policy spectrum on the best approach to stabilize the cycle.

The fact that the evolution of the global neutral rate remains highly relevant for policy and that our analysis has only briefly touched on a vast territory means there are many exciting areas for future research. One extension to the analysis in this paper would be to use a regional perspective to shed light on the fall in global real rates. As discussed in more detail in Rachel and Smith (2015), developed and emerging economies have exhibited different trends in saving and investment over time—future work could usefully explore the reasons for these differences to pinpoint the drivers of global real rates more precisely. Future analysis could also acknowledge the central role of the monetary and financial sector in the global economy of today, by looking at monetary trends that could have affected the global real rate, such as a structural change in global liquidity. This could complement the analysis of real trends highlighted in this paper. Further work could also look to extend our analysis further back in history, in an attempt to explain not only why real rates have fallen since the 1980s but also why they rose in the decades before the 1980s. Finally, development of structural models capable of handling secular, medium- to long-run trends may shed further light on the size of the effects highlighted here, and help clarify how the various effects interact—allowing more in-depth policy experiments and welfare analysis. To the extent that our headline prediction of a persistently low long-run neutral rate comes to pass, such research will remain highly relevant for years to come.

References

Baker, D., J. D. Delong, and P. R. Krugman. 2005. "Asset Returns and Economic Growth." *Brookings Papers on Economics Activity* 2005 (1): 289–330.

Bernanke, B. S. 2005. "The Global Saving Glut and the U.S. Current Account Deficit." Remarks at the Sandridge Lecture, Virginia Association of Economists, Richmond, Virginia, March 10.

———. 2015a. "Why Are Interest Rates So Low?" Brookings Blog Entry, April 1.

———. 2015b. "Why Are Interest Rates So Low? Part 3: The Global Saving Glut." Brookings Blog Entry, March 30.

Bomfim, A. 1997. "The Equilibrium Fed Funds Rate and the Indicator Properties of Term-Structure Spreads." *Economic Inquiry* 35 (4): 830–46.

Broadbent, B. 2014. "Monetary Policy, Asset Prices and Distribution." Speech given at the Society of Business Economists Annual Conference, October 23.

Brynjolfsson, E., and A. McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in Times of Brilliant Technologies*. W.W. Norton & Company.

Buiter, W., E. Rahbari, and J. Seydl. 2015. “Preventing Secular Stagnation from Becoming the ‘New’ Normal: Why the Time of One-Armed Economic Policy Is Gone.” Citigroup Research.

Byrne, D M., S. D. Oliner, and D. E. Sichel. 2013. “Is the Information Technology Revolution Over?” Finance and Economics Discussion Series No. 2013-36, Board of Governors of the Federal Reserve System.

Caballero, R. J., and E. Farhi. 2013. “A Model of the Safe Asset Mechanism (SAM): Safety Traps and Economic Policy.” NBER Working Paper No. 18737.

Carvalho, C., A. Ferrero, and F. Nechio. 2016. “Demographics and Real Interest Rates: Inspecting the Mechanism.” Working Paper No. 5, Federal Reserve Bank of San Francisco.

Christiano, L. J., M. Eichenbaum, and C. L. Evans. 1999. “Monetary Policy Shocks: What Have We Learned and to What End?” *Handbook of Macroeconomics*, Vol. 1, Part A, 65–148 (chapter 2). Elsevier.

Crossley, T., and C. O’Dea. 2010. “The Wealth and Saving of UK Families on the Eve of the Crisis.” Report, Institute for Fiscal Studies.

Cynamon, B. Z., and S. M. Fazzari. 2016. “Inequality, the Great Recession, and Slow Recovery.” *Cambridge Journal of Economics* 40 (2): 373–99.

DeFina, R. H. 1984. “The Link between Savings and Interest Rates: A Key Element in the Tax Policy Debate.” *Business Review* (Federal Reserve Bank of Philadelphia) (November/December): 15–21.

Dynan, K. E., J. Skinner, and S. P. Zeldes. 2004. “Do the Rich Save More?” *Journal of Political Economy* 112 (2): 397–444.

Eggertsson, G. B., and N. R. Mehrotra. 2014. “A Model of Secular Stagnation.” NBER Working Paper No. 20574.

Eichengreen, B. 2015. “Secular Stagnation: The Long View.” NBER Working Paper No. 20836.

Ellis, C., and S. Price. 2003. “UK Business Investment: Long-Run Elasticities and Short-Run Dynamics.” Bank of England Working Paper No. 196.

Fischer, S. 2016. "Why Are Interest Rates So Low? Causes and Implications." Speech at the Economic Club of New York, New York City, October 17.

Frey, C. B., and M. A. Osborne. 2013. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" Paper presented at Oxford University's Machines and Employment Workshop, March 6.

Gilchrist, S., and E. Zakrajsek. 2007. "Investment and the Cost of Capital: New Evidence from the Corporate Bond Market." NBER Working Paper No. 13174.

Goldin, C., and L. F. Katz. 2007. "The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005." NBER Working Paper No. 12984.

Gordon, R. J. 1990. *The Measurement of Durable Goods Prices.* University of Chicago Press and National Bureau of Economic Research.

———. 2014a. "The Demise of US Economic Growth: Restatement, Rebuttal and Reflections." NBER Working Paper No. 19895.

———. 2014b. "US Economic Growth Is Over: The Short Run Meets the Long Run." In *Think Tank 20: Growth, Convergence and Income Distribution: The Road from the Brisbane G-20 Summit*, 185–92. Washington, DC: Brookings Institution.

Guiso, L., A. K. Kashyap, F. Panetta, and D. Terlizzese. 2002. "How Interest Sensitive is Investment? Very (when the data are well measured)." Conference Paper, European Central Bank.

Gutiérrez, G., and T. Philippon. 2016. "Investment-less Growth: An Empirical Investigation." NBER Working Paper No. 22897.

Hamilton, J. D., E. S. Harris, J. Hatzius, and K. D. West. 2015. "The Equilibrium Real Funds Rate: Past, Present and Future." Paper written for the U.S. Monetary Policy Forum, New York City, February 27.

Hansen, A. 1934. "Capital Goods and the Restoration of Purchasing Power." *Proceedings of the Academy of Political Science* 16 (1): 11–19.

———. 1938. "Economic Progress and Declining Population Growth." Presidential address delivered at the 51st annual meeting of the American Economic Association, Detroit, Michigan, December 28.

Havranek, T., R. Horvath, Z. Irsova, and M. Rusnak. 2015. "Cross-Country Heterogeneity in Intertemporal Substitution." *Journal of International Economics* 96 (1): 100–118.

International Monetary Fund. 2014a. "Is It Time for an Infrastructure Push? The Macroeconomic Effects of Public Investment." In *World Economic Outlook: Legacies, Clouds, Uncertainties* (October), 91–112 (chapter 3). Washington, DC: International Monetary Fund.

———. 2014b. "Perspectives on Global Real Interest Rates." In *World Economic Outlook: Recovery Strengthens, Remains Uneven* (April), 75–114 (chapter 3). Washington, DC: International Monetary Fund.

Jorgenson, D. W., and K. M. Vu. 2010. "Potential Growth of the World Economy." *Journal of Policy Modeling* 32 (5): 615–31.

King, M., and D. Low. 2014. "Measuring the 'World' Real Interest Rate." NBER Working Paper No. 19887.

Kumhof, M., R. Ranciere, and P. Winant. 2015. "Inequality, Leverage, and Crises." *American Economic Review* 105 (3): 1217–45.

Kurzweil, R. 2005. *The Singularity Is Near: When Humans Transcend Biology.* The Penguin Group.

Kuznets, S. 1960. "Population Change and Aggregate Output." In *Demographic and Economic Change in Developed Countries*, 324–39. Columbia University Press and NBER.

Laubach, T., and J. C. Williams. 2003. "Measuring the Natural Rate of Interest." *Review of Economics and Statistics* 85 (4): 1063–70.

Lee, R., and A. Mason, eds. 2011. *Population Aging and the Generational Economy: Global Perspective.* Edward Elgar.

Mendes, R. R. 2014. "The Neutral Rate of Interest in Canada." Discussion Paper No. 2014-5, Bank of Canada.

Mishkin, F. S. 1984. "Are Real Interest Rates Equal Across Countries? An Empirical Investigation of International Parity Conditions." *Journal of Finance* 39 (5): 1345–57.

Piketty, T. 2014. *Capital in the Twenty-First Century.* Harvard University Press.

Piketty, T., and E. Saez. 2014. "Inequality in the Long Run." *Science* 344 (6186): 838–43.

Rachel, L., and T. D. Smith. 2015. "Secular Drivers of the Global Real Interest Rates." Staff Working Paper No. 571, Bank of England.

Ramey, V. A. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics,* Vol. 2, ed. J. B. Taylor and H. Uhlig, 71–162 (chapter 2). Elsevier.

Saez, E., and G. Zucman. 2014. "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data." NBER Working Paper No. 20625.

Speller, W., G. Thwaites, and M. Wright. 2011. "The Future of International Capital Flows." Financial Stability Paper No. 12, Bank of England.

Summers, L. H. 2013. "Economic Possibilities for our Children." Remarks delivered at the Fifth Annual Martin Feldstein Lecture, held at the NBER Summer Institute, Cambridge, Massachusetts, July 24. Published in the *NBER Reporter* 2013 No. 4.

———. 2014a. "Bold Reform Is the Only Answer to Secular Stagnation." *Financial Times*, September 7.

———. 2014b. "U.S. Economic Prospects: Secular Stagnation, Hysteresis and the Zero Lower Bound." *Business Economics* 49 (2): 65–73.

Thwaites, G. 2015. "Why Are Real Interest Rates So Low? Secular Stagnation and the Relative Price of Investment Goods." Discussion Paper, Centre for Macroeconomics.

Discussion of "Are Low Real Interest Rates Here to Stay?"*

Pierre-Olivier Gourinchas
University of California at Berkeley

1. The Trillion Dollar Question

The decline in global real interest rates since the early 1980s is nothing short of spectacular. In the United States, the entire yield curve on government securities has shifted down by more than 10 percent between 1983 and 2013. A similar pattern, although perhaps not as pronounced, has been documented in most advanced and many emerging economies. According to the authors, long-term global real interest rates have declined by more than 450 basis points (bps) since the early 1990s. Coupled with low and stable worldwide inflation, this overwhelmingly suggests a decline in the global "neutral" or "natural" rate of interest, colloquially known as r^*, of a similar magnitude.[1]

It is difficult to overstate the macroeconomic importance of such a phenomenon. Among other things, a decline in r^* of this magnitude seriously limits the ability of monetary authorities to stabilize aggregate demand, as policy rates cannot be lowered below the "effective" lower bound. It also poses specific challenges to financial stability. For instance, excessive cuts in the policy rate can lower banks' net interest margin, with adverse effects on the health of the banking sector. A low-r^* world also creates powerful incentives for financial intermediaries to engineer "pseudo" safe assets, and for investors to load up excessively on higher-risk/higher-yield assets.

*This discussion was prepared for the November 2016 IJCB conference held at the Federal Reserve Bank of San Francisco. Part of this discussion builds upon joint work with Hélène Rey.

[1]The paper presents a useful discussion of the subtle differences between the global neutral rate and country-specific r^*, but these subtleties are not relevant for the purpose of this discussion since the purpose of the paper is on the determinants of the global real rate, and the authors argue that this global neutral rate is a long-run anchor for r^*.

The trillion-dollar question, then, is to understand the sources of this decline. Not surprisingly, given the importance of the topic, potential explanations abound—from a decline in long-run productivity growth to demographic forces due to an aging global population; from rising income inequality to a decline in the relative price of investment goods; or from increased foreign exchange reserve accumulation by emerging market central banks to a shift in investors' desired portfolio weights towards safe assets, among others—so much so that one might be tempted to conclude that the decline in global real rates is, if anything, *over-explained*.

Sifting carefully through these different explanations and sorting out which are relevant and which are not is no easy task. Yet it is the one that the authors, somewhat heroically in my view, set for themselves in this paper. To do so, they propose a very detailed and nuanced quantitative assessment of the importance of each of the large number of factors mentioned above, and some others.

The final decomposition looks both sensible and interesting: of the 450 bps decline in long-term real yields, the authors claim to account for 400 bps—100 bps due to a decline in future trend growth, 160 bps due to increased desired savings, and 140 bps from decreased desired investment. Two determinants appear particularly relevant in the authors' decomposition: the impact of demographic forces on desired savings (90 bps) and that of a rising spread between risk-free rates and the return to capital, dampening desired investment (70 bps).

I am very sympathetic to that final point, in part because it is very much in line with my own work on the topic! For instance, in Caballero, Farhi, and Gourinchas (2017), we propose a decomposition of that same spread into a risk premium component, and a component due to increased markups ("rents") or to increased capital-augmenting technological progress ("automation"). That paper concluded that while increased rents or automation cannot be ruled out, they alone cannot account for the increased spread between returns to capital and safe real interest rates. Instead, a substantial part of the increase in that spread reflects increased compensation for risk, indicating a substantial increase in the demand for, or a significant decline in the supply of, safe assets relative to risky assets.

However, while my general views are very sympathetic with the general conclusions of the paper, I am much less convinced by the methodology that the authors have adopted in their paper. The next section reviews my main concerns in that respect. Next, I present an altogether different approach, based on the intertemporal budget constraint, and my recent work with Hélène Rey (Gourinchas and Rey 2017). That framework, using long-run data, finds also an important role for the financial cycle, with real interest rates remaining low for an extended period of time following global financial crises.

2. Accounting for Changes in the Natural Interest Rate

At its core, an economic model is a mapping from shocks $\bar{\varepsilon} = \{\varepsilon_i\}$ to observables $Y = \{y_j\}$. That mapping is controlled by a number of structural parameters Θ and can be summarized as $Y = f(\bar{\varepsilon}; \Theta)$.[2]

The first step consists in estimating the fundamental parameters $\hat{\Theta}$. An impulse response can then be obtained by setting some component i of ε to 1 and all other shocks to 0. An accounting decomposition can be obtained once the fundamental shocks $\hat{\bar{\varepsilon}} = \{\hat{\varepsilon}_i\}$ are also estimated. The contribution of shock i can then be obtained by constructing $Y_i = f(\hat{\bar{\varepsilon}}_i; \bar{\Theta})$, where $\hat{\bar{\varepsilon}}_i = \{\hat{\varepsilon}_i, 0\}$ ignores all other shocks beside shock i.[3]

While the general principles are well understood, their application is often difficult. For instance, proper identification of the structural parameters Θ is often difficult to achieve. Precise estimation of the structural shocks can also be difficult.

In the context of this paper, such a decomposition exercise would start with a model incorporating all the potential sources of decline in the neutral rate of interest listed above: productivity growth, demographic forces, income inequality, the relative price of investment goods, foreign exchange reserve accumulation, etc. Notice that some of these sources, such as income inequality or the relative price of investment goods, are themselves endogenous to, inter alia, the structure of the tax system; the degree of competition in factor

---

[2]That definition is quite generic and does not impose that the model be static: y_j can include current, lagged, and even future observations of macroeconomic variable j, and similarly ε_i can include current, lagged, and future values of shock i.

[3]With non-linearities, these decompositions do not necessarily add up to 100 percent since the full response also incorporates interaction terms.

markets; productivity in the human accumulation sector; unobserved abilities; or productivity in the investment good sector, etc. The full model would incorporate a large number of potential shocks and a large number of structural parameters. Unfortunately, the complexities and identification challenges of estimating such a model lie well beyond our technical abilities.

Instead, this paper cobbles together estimates from two different approaches. The first approach is a Euler equation approach:

$$r = \frac{1}{\sigma} g + \rho, \tag{1}$$

where r is the real interest rate, g is the growth rate of technology, σ is the intertemporal elasticity of substitution (IES), and ρ is the rate of time preference. This equation obtains from the optimal intertemporal allocation of consumption by households. The interpretation is straightforward: faster productivity growth means more output (and consumption) tomorrow relative to today. Households want to smooth consumption over time, hence they would like to borrow and consume more today. This pushes up the real interest rate, the more so the lower is the IES (low σ). For a given IES, this maps changes in g into changes in r.

The second approach is an $S - I$ diagram. It consists of identifying shifts in desired savings S and desired investment I arising from the different mechanisms mentioned above. For given estimates of the interest rate elasticity of the saving and investment curves, it maps the changes in desired savings and investment into changes in the equilibrium real interest rate.

Unfortunately, while this approach is more tractable than the full-blown estimation of a structural model, it leads to a number of conceptual difficulties.

To start with, many of the "channels" discussed can have complex effects on the real interest rate, so the issue of identification remains unsettled. Consider for instance the link between the dependency ratio (defined as the fraction of the population not of working age) and savings. The empirical evidence documented in the paper indicates a weak negative relationship between the saving rates and dependency ratios in the cross-section, which is stable over time (figure 6). The authors assume that this relationship also holds in the aggregate time series. Since the dependency ratio has declined globally, they conclude that this contributed to increased desired aggregate savings. Given an assumed elasticity of investment of -0.7, this accounts for 90 bps of the 160 bps declined in interest rates due to increased savings. This is one of the largest single contributors to the decline in real interest rates. If true, this would suggest that future increases in the dependency ratio would *raise* real interest rates as aggregate savings decline. Yet, it is well known that some countries with a relatively high dependency ratio due to rapid population aging, such as Japan or Germany, are also among the countries with the highest saving rate.

The answer may well lie in the fact that the dependency ratio can vary when the proportion of either young dependents or old dependents varies. But the implications for aggregate savings may be vastly different. Historically, countries with low dependency ratios are countries that successfully made their demographic transition. Fewer children also mean higher savings, since in many developing countries children were effectively a means of saving for old age. As an illustration, China's one-child policy may well have contributed to an increase in Chinese savings. As the proportion of young dependents has decreased globally, savings have increased. But the implications may well be very different if the change in the dependency ratio comes from the fraction of old dependents. An increase in population aging means more retirees relative to the working-age population. It is likely to be associated with an *increase*, not a decrease, in savings. Once this heterogeneity is taken into account, it is not so clear that the global dependency ratio is a good summary statistic for the desired savings shift due to demographic forces. The authors' conclusion from this reduced-form exercise seems relatively fragile.

Assuming that desired shifts in savings and investment are properly measured, translating these into shifts in equilibrium real interest rates requires reliable estimates of the savings and investment interest rate elasticities. The authors assume an elasticity of savings of 0.5, but mention a range from 0.03 to 1.8. From a theoretical and empirical point of view, it is well known that the relationship between savings and interest rates can be quite fickle. It is in fact quite telling that the most recent paper cited by the authors dates back to 1983: the literature seems to have concluded that the elasticity of aggregate savings to the real interest rate is not a well-defined parameter and one should perhaps avoid trying to estimate it. Yet it is a critical component of the exercise. The interest rate elasticity of investment is also imprecisely estimated, between –0.5 and –1, although at least both the theoretical and empirical literature agree that it should be negative.

Finally, it is not clear how the $S-I$ framework is an alternative to the Euler equation approach. After all, the Ramsey-Cass-Koopman model embeds an $S-I$ diagram too. A decline in productivity growth lowers the marginal product of capital, which reduces the desired capital stock and reduces investment. In equilibrium the real interest rate needs to decline to reduce saving, so that $S = I$ obtains. In other words, the $S-I$ diagram is precisely what delivers $r = g/\sigma+\rho$. It is not clear to me whether we are not double-counting when we are considering the effect of, e.g., demographics or rising inequality on trend productivity via the Euler equation and the $S - I$ diagram. One of the advantages of the structural model $Y = f(\bar{\varepsilon}, \Theta)$ is precisely that it avoids such a double-counting.

Finally, the Euler equation is a weak peg to hang the empirical estimates on. A large and abundant empirical literature has documented very weak support for the aggregate Euler equation and the absence of a strong relationship between real interest rates and growth.

The bottom line is that, while the structural approach may be infeasible, the approach followed in this paper, cobbling and adding together many reduced-form estimates of "desired" shifts in savings and investment, combined with an Euler equation, does not provide solid empirical estimates of the relative contributions of the various channels. Unfortunately, no matter how complex the problem is, we cannot afford to avoid using some structural approach. In doing so, the trade-off is between parsimony and tractability. The next section presents the results from such a framework, borrowed from Gourinchas and Rey (2017).

3. An Alternative Framework to Understand Secular Movements in r

In Gourinchas and Rey (2017), we propose an accounting framework based on the movements in the ratio of total consumption expenditures C to total wealth W over long periods of time. Figure 1 plots this ratio for a world that comprises the United States, the United

Figure 1. The Global Consumption-Wealth Ratio

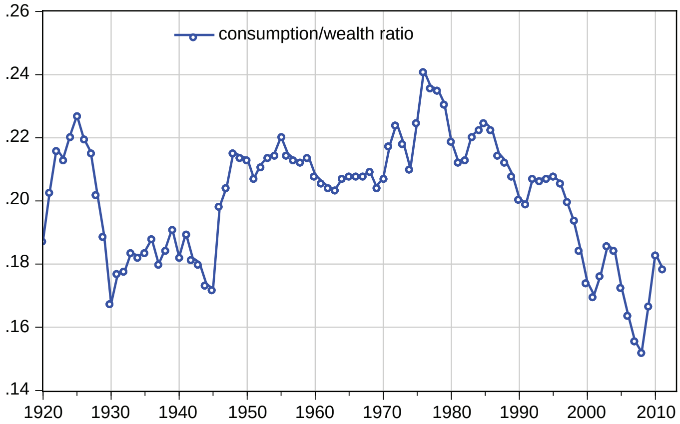

Source: Jordà, Schularick, and Taylor (2016), Piketty and Zucman (2014).

Note: The figure reports the ratio of aggregate annual private consumption expenditures to total private wealth (land, housing, financial assets) for the United States, the United Kingdom, Germany, and France.

Kingdom, Germany, and France between 1920 and 2011. Wealth consists of total private wealth, including land, housing, and financial assets as estimated by Piketty and Zucman (2014). The ratio, which can be interpreted as an average propensity to consume out of wealth, exhibits low-frequency fluctuation, between a low of 0.155 in 2009 and a high of 0.24 in 1975.

Under the mild assumption that this ratio is stationary, an assumption consistent with most theories of consumption, we can write the following decomposition:

$$\ln(C_t/W_t) \equiv cw_t = \sum_{s=0}^{\infty} \rho^s \mathbb{E}_t r_{t+s} + \sum_{s=0}^{\infty} \rho^s \mathbb{E}_t rp_{t+s} - \sum_{s=0}^{\infty} \rho^s \mathbb{E}_t g^C_{t+s}$$
$$= cw_t^{rf} + cw_t^{rp} + cw_t^c. \qquad (2)$$

In this equation, the log consumption-to-wealth ratio cw_t is the sum of three components: the present discounted sum of future risk-free rates r_{t+s}, the present discounted value of future risk premia

Figure 2. Decomposing *cw*

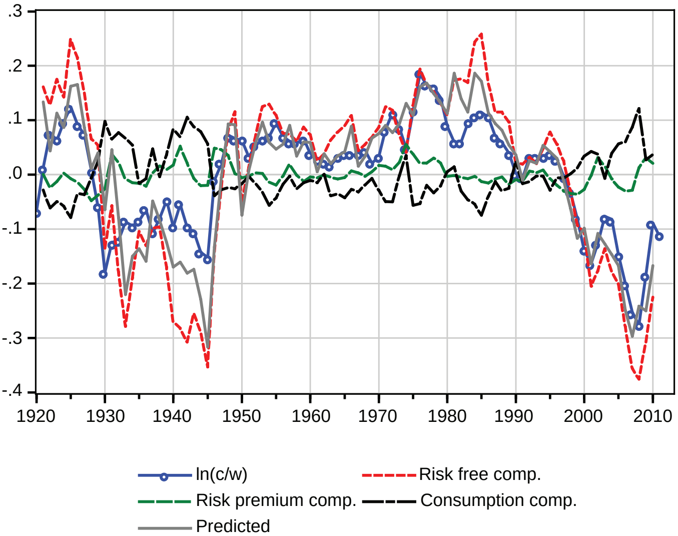

Note: The figure decomposes $\ln(C/W)$ into a risk-free component (cw^{rf}), an excess return component (cw^{rp}), and a consumption growth component (cw^{c}).

rp_{t+s}, and the present discounted sum of future aggregate consumption growth rates g^C_{t+s}. The discount rate ρ is a constant that depends on the steady-state value of C/W. The interpretation of this formula is quite straightforward: if C/W is stationary, a low C/W value today predicts a higher value in the future. The increase in C/W must come from either an increase in the numerator, i.e., high future consumption growth g^C_{t+s}, or a decline in the denominator, i.e., low returns on wealth, which implies either low future risk-free rates r_{t+s} or low future risk premia rp_{t+s} or both.

Using a vector autoregression approach, Gourinchas and Rey (2017) estimate the three components on the right-hand side of this equation. Figure 2 reports the estimated components, together with the overall fit of the regression. Because equation (2) is the approximation of an accounting identity (the global budget constraint), we expect the overall fit to be high, as is indeed the case. More importantly, the figure reveals that the risk-free component cw^{rf} accounts for most of the variations in cw while other components (future consumption growth or future risk premia) play relatively minor roles. In other words, a low consumption-to-wealth ratio today indicates that the present discounted value of future risk-free rates is below average. Over long periods of time, because of the effective lower bound, the real interest rate is an upper bound on the natural rate r^*. It follows that a low value of cw is associated with low future r^*. In Gourinchas and Rey (2017), we discuss how this result suggests that the slowdown in productivity growth or demographic forces are unlikely to be the main drivers of the decline in r^*. The basic insight is that low future productivity growth or a slowdown in population growth would both reduce the future growth rate of aggregate consumption growth g^C_{t+s}. According to equation (2), this would tend to *raise* cw today, not depress it, unless the equilibrium response of real interest rates to a growth slowdown, i.e., the intertemporal elasticity of substitution $1/\sigma$ in equation (1), is sufficiently low. In that case, cw^{rf} and cw^c would be negatively correlated and their ratio, an estimate of σ, would be implausibly low.

The upshot of the exercise is that either we live in a world of very low IES, so that global interest rates are extremely responsive to changes in productivity growth (more than one for one), or the movements in r originate elsewhere. Our favorite hypothesis is that of financial boom-bust cycles casting a long shadow on future real rates. In particular, the two episodes of persistently low cw seem to occur in the run-up to and the aftermath of global financial cycles: a Great Depression episode from from 1925 to 1945, and the global financial crisis from 2000 onwards. In both cases, the decline in cw arises first from a rapid run-up in asset and housing prices that increases wealth (the denominator) that is not matched by a corresponding increase in consumption. The financial crisis, when it occurs, should have a corrective effect, restoring cw back to its trend. However, we observe instead that cw remains depressed. One explanation is that the financial collapse triggers increased demand for safe assets and economy-wide deleveraging dynamics: as households, corporates, and governments all simultaneously try to reduce their borrowing levels, safe interest rates collapse. Our findings also suggest that though risk premia seem to increase, as expected if the demand for safe assets increases, this is not sufficient to offset the decline in risk-free rates, and the overall expected return on wealth remains depressed post-crisis.

Figure 3. Predicting Global Risk-Free Rates

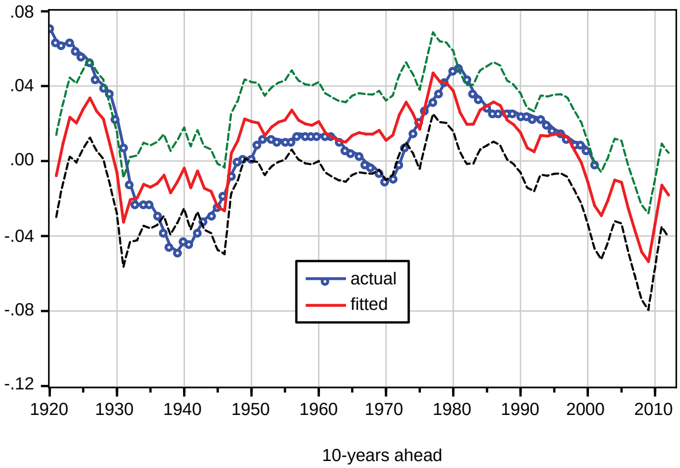

Note: The figure forecasts the ten-year average future short risk-free rate using $\ln(C/W)$. The graph includes two standard deviation bands.

Finally, our decomposition allows us to answer directly the question posed by the authors: Are low real interest rates here to stay? Figure 3 reports the predicted values based on a regression of a ten-year average of the real risk-free rate on the initial value of cw. The line with circles reports this average future interest rate between 1920 and 2001 (since our data cover the period 1920–2011). The solid line reports the estimated future risk-free interest rate, with two standard deviation confidence bands. The in-sample fit is very high. Extrapolating past 2001 until 2011, it reveals that global real risk-free rates are expected to remain low for an extended period of time. As of the last data point in our sample, 2011, the average short-term real risk-free rate over the subsequent ten years, 2011–21, is expected to be –2 percent.

4. Conclusion

The paper presents an ambitious attempt to account quantitatively for the different economic forces behind the recent decline in real interest rates. This is a difficult empirical exercise. Without a formal structural framework, identification is weak, and it is difficult to isolate cleanly the effect of various forces. Yet the exercise remains useful as a starting point for more structural explorations. I have presented the results from one such exercise, based on Gourinchas and Rey (2017). It provides a complementary set of results and emphasizes the historical role of deleveraging dynamics in depressing real interest rates.

References

Caballero, R. J., E. Farhi, and P.-O. Gourinchas. 2017. "Rents, Technical Change, and Risk Premia Accounting for Secular Trends in Interest Rates, Returns on Capital, Earning Yields, and Factor Shares." *American Economic Review* 107 (5): 614–20.

Gourinchas, P.-O., and H. Rey. 2017. "Real Interest Rates, Imbalances and the Curse of Regional Safe Asset Providers at the Zero Lower Bound." In *The Future of the International Monetary and Financial Architecture*. Proceedings of the ECB Forum on Central Banking, sponsored by the European Central Bank, held in Sintra, Portugal, June 27–29, 2016.

Jordà, Ò., M. Schularick, and A. M. Taylor. 2017. "Macrofinancial History and the New Business Cycle Facts." *NBER Macroeconomics Annual 2016*, ed. M. Eichenbaum and J. A. Parker, 213–63. National Bureau of Economic Research.

Piketty, T., and G. Zucman. 2014. "Wealth and Inheritance in the Long Run." CEPR Discussion Paper No. 10072 (July).

Monetary Policy, the Financial Cycle, and Ultra-Low Interest Rates*

Mikael Juselius,[a] Claudio Borio,[b] Piti Disyatat,[c] and Mathias Drehmann[b]

[a]Bank of Finland

[b]Bank for International Settlements

[c]Bank of Thailand

Do the prevailing unusually and persistently low real interest rates reflect a decline in the natural rate of interest as commonly thought? We argue that this is only part of the story. The critical role of financial factors in influencing medium-term economic fluctuations must also be taken into account. Doing so for the United States yields estimates of the natural rate that are higher and, at least since 2000, decline by less. An illustrative counterfactual experiment suggests that a monetary policy rule that takes financial developments *systematically* into account during both good and bad times could help dampen the financial cycle, leading to significant output gains and little change in inflation.

JEL Codes: E32, E40, E44, E50, E52.

1. Introduction

Inflation-adjusted (real) interest rates, short and long, have been on a downward trend for a long time and have remained exceptionally low since the Great Financial Crisis (figure 1). Why is this so?

*We would like to thank Raphael Auer, Stijn Claessens, Marc Giannoni, Boris Hofmann, Jonathan Kearns, Giovanni Lombardo, Phurichai Rungcharoenkitkul, Anders Vredin, Hyun Shin, and John Williams for helpful comments and discussions, and Andreas Freitag for excellent statistical assistance. All remaining errors are ours. The views expressed are those of the authors and do not necessarily represent those of the Bank for International Settlements, the Bank of Finland, or the Bank of Thailand. Author e-mails: Juselius (corresponding author): mikael.juselius@bof.fi; Borio: claudio.borio@bis.org; Disyatat: pitid@bot.or.th; Drehmann: mathias.drehmann@bis.org.

Figure 1. The Long-Term Decline in Real Interest Rates

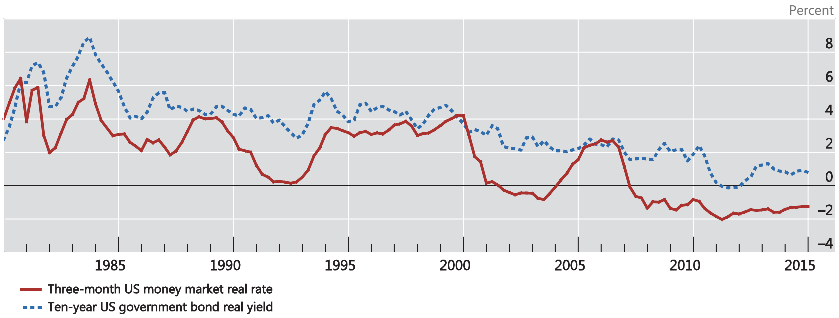

Source: National data.

Note: Real rates are generated by subtracting realized PCE core inflation from nominal interest rates.

The prevailing view is that this downward trend and the exceptionally low level largely reflect a fall in natural interest rates, driven by changes in saving and investment fundamentals (International Monetary Fund 2014, Bean et al. 2015, Obstfeld and Tesar 2015, Rachel and Smith 2015). One prominent variant is the hypothesis that persistently weak demand for capital, a rising propensity to save, and lower trend growth have brought about an era of "secular stagnation" (Summers 2014). Another variant points to a higher propensity to save in emerging economies together with investors' growing preference for safe assets (Caballero, Farhi, and Gourinchas 2008, Broadbent 2014, Bernanke 2015, Del Negro et al. 2017).

Views about the natural rate are necessarily model dependent. At the heart of the prevailing interpretation are two key features. First, the natural rate is defined as that which would prevail when actual output equals potential output. Second, inflation is the key signal that output is not at its potential, sustainable, level. All else equal, if output is above potential, inflation will tend to rise; if it is below, inflation will tend to fall. The natural rate, so defined, is also known as the "Wicksellian" rate, following Wicksell (1898) and as refined by Woodford (2003). Crucially, this view presumes that over the medium term, monetary policy only passively tracks the natural rate. As a result, the observed decline in real interest rates is purely a function of forces beyond central banks' control.

We argue that this view is too narrow. When we think of both potential output and the symptoms of unsustainability, we also need to consider financial factors. Output cannot be at a sustainable level if the financial side of the economy is out of kilter. And the key symptom of unsustainability may be outsized financial booms and busts, which in turn can wreak havoc on output (e.g., Borio and Lowe 2002, Schularick and Taylor 2012). Indeed, empirical evidence indicates that many recessions, especially those that coincide with banking crises, have permanent effects on output—growth may return to its pre-crisis long-term trend, but output does not, so that a permanent gap develops between the pre- and post-recession output trends (Cerra and Saxena 2008, Basel Committee on Banking Supervision 2010, Ball 2014).

We present an alternative view of the natural rate, in which financial factors also play a role. This has a couple of advantages. Analytically, it avoids the conclusion that interest rates may be at their long-run equilibrium or natural level and yet encourage the buildup of serious financial instability (e.g., Summers 2014, Bean et al. 2015). Empirically, our view dovetails with the burgeoning literature documenting the limited usefulness of domestic measures of slack (deviations of output from potential) as determinants of inflation (e.g., Borio and Filardo 2007, Pain, Koske, and Sollie 2008, Ball and Mazumder 2011, IMF 2014, Blanchard, Cerutti, and Summers 2015)[1] and—the mirror image—with the limited usefulness of inflation as an indicator of business-cycle conditions (e.g., Borio, Disyatat, and Juselius 2014). Indeed, a recent strand of empirical work indicates that, by contrast, financial-cycle proxies are helpful indicators of those conditions (Borio, Disyatat, and Juselius 2014, 2017; Kiley 2015).

This perspective has first-order implications for monetary policy. It suggests that inflation may be an insufficient guide for monetary policy. If monetary policy has a material impact on financial booms and busts, thereby resulting in swings in medium-term

[1] Consistent with this finding is the recent evidence that domestic output gaps in standard Phillips-curve models provide little additional predictive content beyond lagged inflation when forecasting inflation (Dotsey, Fujita, and Stark 2015). Faust and Leeper (2015) also stress that inflation dynamics are not as simple as implied by models based on economic slack.

output trajectories, and if inflation is a poor indicator of deviations of output from potential, then ignoring financial cycles may lead policy astray. To analyze these issues in more detail, we propose an empirical framework in which financial factors play a pivotal role in economic fluctuations.[2] Our objective is twofold: (i) to revisit the measurement of the natural interest rate; and, more ambitiously, (ii) to propose a monetary policy rule that *systematically* takes into account the state of the financial cycle. By establishing a link between monetary policy and the financial cycle, the framework also provides a richer perspective on the secular decline in real interest rates.

We apply the framework to U.S. data over a thirty-year period, 1985–2015, and reach two main conclusions. First, once financial factors are taken into account, the natural interest rate is higher and falls by less than prevailing empirical approaches would suggest, at least since 2000. Importantly, the actual real policy interest rate has been persistently below the natural rate, especially in the most recent period. Second, the way monetary policy is systematically conducted has a first-order impact on financial factors and hence output fluctuations. And the resulting booms and busts have persistent effects even when crises do not break out. As a corollary, taking the financial cycle into account can improve economic performance.

Together, these two conclusions suggest that a narrative that attributes the decline in real interest rates and their persistently ultra-low post-crisis levels to an exogenous fall in the natural rate is incomplete. Monetary policy, through its impact on the financial cycle, influences the evolution of real interest rates over the medium term. In this sense, beyond the structural evolution of the

[2]There is, of course, a much broader literature highlighting the role of financial factors in economic fluctuations and, more specifically, the GDP costs of financial or credit booms and busts. Some of those closest in spirit to the analysis performed here include, for instance, Claessens, Kose, and Terrones (2009), Reinhart and Rogoff (2009), Schularick and Taylor (2012), Jordà, Schularick, and Taylor (2013), and Mian, Sufi, and Verner (2016). On the theoretical side, following Bernanke, Gertler, and Gilchrist (1999), a prominent strand of research has focused on the role of financial frictions in New Keynesian models with increasing emphasis on variations in credit spreads (e.g., Del Negro, Giannoni, and Schorfheide 2015, Adrian and Duarte 2016, Cúrdia and Woodford 2016, Del Negro et al. 2017).

economy, the decline also reflects, in part, policy frameworks (Borio and Disyatat 2014, Borio 2016).

Thus, monetary policy frameworks matter. An effective "lean against the wind" approach requires policy to take financial developments into account *systematically.* One possible representation of such a policy rule, as described here, takes the form of an augmented version of the standard Taylor rule (Taylor 1993) that incorporates financial-cycle indicators. Such a rule differs fundamentally from typical interpretations of a "lean against the wind" policy, whereby interest rates are raised only when signs of financial imbalances, such as credit and asset price run-ups beyond historical norms, emerge (e.g., Svensson 2014, 2016; Ajello et al. 2015; IMF 2015). Responding to financial stability risks only when they become evident would lead to doing too little too late, as it would ignore the *cumulative* impact of policy over the whole financial cycle. Rather, policy interest rates should be set so that the economy is never too far away from "financial equilibrium"—a notion that we will define more precisely below. Using an illustrative policy rule that embodies such features, our analysis suggests that it would have been possible to mitigate financial imbalances, leading to significant output gains.

To reach these conclusions, we expand the familiar Laubach and Williams (2003, 2015a) reduced-form model for estimating potential output and the natural rate by incorporating financial-cycle information. Our measure of the financial cycle focuses on the role of leverage and debt service burdens, assessed relative to their long-run levels. As Juselius and Drehmann (2015) find, such a measure can account for a substantial part of the evolution of output in the short to medium term.[3] By harnessing financial-cycle information, we obtain different estimates of potential output and the natural real interest rate—what we refer to as *"finance-neutral"* estimates. Intentionally, we make the smallest possible adjustments to the system put forward by Laubach and Williams, which relies heavily on the information content of inflation. By nesting this standard framework in ours, we let the data speak. Hence, if the financial terms are

[3]In particular, the dynamics implied by these two gaps captures well and out of sample the basic features of the Great Recession and of the subsequent weak recovery (Juselius and Drehmann 2015).

unimportant, the estimates will simply revert to more conventional estimates.

Based on this framework, we then explore the interaction between policy and the financial cycle. We use the estimates of the filter and perform a counterfactual experiment with a policy rule that takes financial factors systematically into account—the augmented Taylor rule noted above. This part of the exercise is necessarily more speculative, as it faces well-known and serious econometric challenges. A key one is the "Lucas critique": there is no presumption that the estimated coefficients are invariant to policy. Unfortunately, the question we wish to address makes such a critique inevitable, since we are interested in the *systematic* part of policy, not in small and, above all, temporary deviations from an established pattern. We draw some reassurance from the possibility that our results would actually be reinforced if agents were to internalize the systematic policy reaction and hence respond even more strongly. In any case, our aim here is simply to illustrate how a systematic interaction between policy and the financial cycle can have persistent effects on economic performance.

The rest of the paper is organized as follows. Section 2 develops extensions to a baseline system in the spirit of Laubach and Williams (2003). Section 3 introduces the data and presents the estimates. Section 4 simulates the system under different specifications of the policy rule. Possible avenues for future work are considered in the conclusion.

2. The Financial Cycle, Potential Output, and the Natural Interest Rate

Before turning to the detailed extension of the small-scale version of the standard Laubach-Williams framework to take into account financial-real linkages, it is useful to introduce our empirical measure of the financial cycle.

2.1 Characterizing the Financial Cycle

To model the financial cycle, we build on previous work by Juselius and Drehmann (2015). The authors use two co-integrating (long-run) relationships to pin down the long-term evolution of three key variables, namely the credit-to-GDP ratio, real asset prices, and the nominal lending rate. They show that deviations from these long-run relationships provide a useful decomposition of the financial cycle.

The first long-run relationship is between the credit-to-GDP ratio and inflation-adjusted (real) asset prices. This relationship captures the well-known positive link between debt and the price of assets, which may arise from the latter's use as collateral or, more generally, as a source of revenue or service streams (housing). The relationship can be interpreted as a very rough proxy for aggregate *leverage* at market prices. As it turns out, this is well approximated by the ratio of debt to assets (Juselius and Drehmann 2015). We refer to the deviations of this relationship from its long-run value as the *leverage gap*, $\widetilde{lev}_t$, which can be expressed as[4]

$$\widetilde{lev}_t = (cr_t - y_t) - (p_{A,t} - p_t) - \overline{lev}, \tag{1}$$

where cr_t is credit to the non-financial private sector, y_t is output, $p_{A,t}$ is a real asset price index, p_t is the consumer price level, and $\overline{lev}$ is a steady-state constant. The asset price index is constructed from residential property prices, commercial property prices, and equity prices, with the corresponding weights adding up to one (see appendix). When leverage is below its long-run value (e.g., when asset prices are high), credit tends to increase faster than normal, and vice versa.

The second long-run relationship is between the credit-to-GDP ratio and the (average) lending rate on the debt outstanding. This relationship captures the link between debt and interest payments, consistent with the notion that a lower interest bill allows households and firms to service the same stock of debt with lower income in the long run. As it turns out, this is a very good proxy for the actual private sector's *debt service ratio*—defined as the ratio of interest payments plus amortizations of households and non-financial companies to their income (Juselius et al. 2016). This suggests that the variable is closely linked to cash flow constraints. We refer to deviations of

[4]In what follows, we use lowercase letters to denote the natural logarithm of a variable—for example, $y_t = \ln(Y_t)$ for the log of nominal GDP—except for the interest rate, which is in levels.

Figure 2. Evolution of the Leverage and Debt Service Gaps

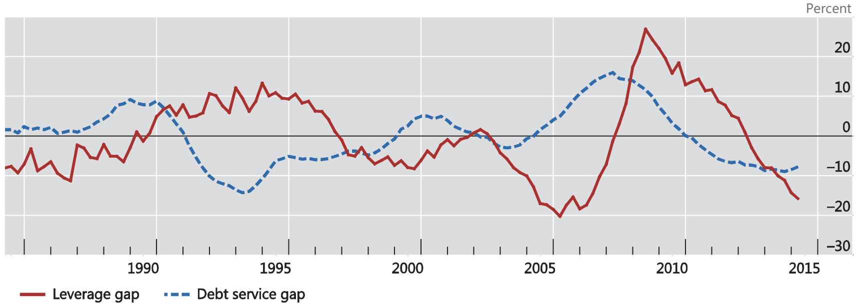

Source: Authors' calculations.

this variable from its long-term relationship as the *debt service gap*, $\widetilde{dsr}_t$, which can be written as

$$\widetilde{dsr}_t = (cr_t - y_t) + \beta_{dsr} i_{L,t} - \overline{dsr}, \tag{2}$$

where $i_{L,t}$ is the nominal average lending rate on the stock of credit, and $\overline{dsr}$ is a steady-state constant. Importantly, $i_{L,t}$ reflects not only *current* interest rate conditions but also past money market rates, past interest rate expectations, and past risk premia, as embedded in the stock of outstanding contracts. Thus, it is influenced by current *and* past monetary policy decisions.

Together, relationships (1) and (2) pin down the long-run level of the credit-to-GDP ratio, consistent with real asset prices (via the leverage gap) and the nominal lending rate (via the debt service gap). In effect, when *both* leverage and debt service gaps are closed, the credit-to-GDP ratio, real asset prices, and the lending rate take values that are consistent with their long-run levels. We adopt this as our measure of *financial equilibrium.*

Figure 2 depicts estimated leverage and debt service gaps for the United States from 1985:Q1 to 2015:Q1.[5] The debt service gap was

[5]Details of the estimation of the co-integration system are set out in the appendix. Commercial property prices are not available before 1985 and, hence, constrain the sample period.

large and positive before and during the three recessions in our sample, in particular for the most recent one. By contrast, the leverage gap was very low during the commercial real estate and leveraged buyout (LBO) boom in the late 1980s and the housing boom in the mid-2000s. This simply reflects the fact that asset prices tend to run ahead of the credit-to-GDP ratio during booms, even as this ratio increases beyond historical trends. In other words, while the credit-to-GDP ratio soars during a credit boom, the leverage gap, as measured here, actually declines, because asset prices increase even more. This also makes borrowers look deceptively solid in the boom phase.

2.2 A "Finance-Neutral" Filtering System

If financial factors play an important role in driving economic fluctuations and contain information about the economy's cyclical variations, then ignoring them is bound to provide less accurate estimates of sustainable trajectories and trends. Building on previous work which has found that financial variables can help improve estimates of potential output (e.g., Arseneau and Kiley 2014, Borio, Disyatat, and Juselius 2014, 2017), we construct a filtering system that allows financial factors to play a role in business-cycle fluctuations and then jointly estimate what might be termed the "finance-neutral" natural interest rate and potential output—in the sense that the estimates control for the influence of financial factors. Specifically, we modify *as little as possible* the standard baseline framework of Laubach and Williams (2003) to take the financial cycle into account.

The baseline system consists of four key equations. First, there is a reduced-form IS equation linking the output gap, $y_t - y_t^*$, to the difference between the real rate and the natural rate, $r_t - r_t^*$ (the interest rate gap), where $r_t = i_t - E_t\pi_{t+1}$. For simplicity, we assume that expected and actual inflation coincide, i.e., $E_t\pi_{t+1} = \pi_t$, throughout.[6] Second, there is a standard Hodrick-Prescott (HP)

[6]This is similar in spirit to Laubach and Williams (2003), who use a weighted average of current and past inflation rates as a proxy for expected inflation. In general, our system has a less elaborate autoregressive structure than theirs, and we link potential output and the natural rate directly rather than adding an additional equation for potential output growth. This leaves small high-frequency components in our potential output and the natural rate estimates, but does not otherwise affect the results.

specification for potential output. The specification is very flexible and can even accommodate a trending growth rate, but we later impose technical restrictions that anchor it to actual output over the medium to long run. Third, there is a Phillips curve, which includes an inflation target, π^*. Finally, there is an equation linking the natural interest rate to the growth rate of potential output and a term, z_t, capturing other determinants of the natural rate, such as the rate of time preference. The system is closed with an equation describing the evolution of z_t.

With these choices, the baseline system, shown in black, becomes

$$y_t - y_t^* = \beta_3(y_{t-1} - y_{t-1}^*) - \varphi_{31}(i_{t-1} - \pi_{t-1} - r_{t-1}^*) - \varphi_{32}\widetilde{lev}_{t-1} + \vartheta_{3t} \tag{3}$$

$$y_t^* = 2y_{t-1}^* - y_{t-2}^* + \vartheta_{4t} \tag{4}$$

$$(\pi_t - \pi^*) = \beta_5(\pi_{t-1} - \pi^*) + \varphi_5(y_{t-1} - y_{t-1}^*) + \vartheta_{5t} \tag{5}$$

$$r_t^* = \beta_6 r_{t-1}^* + (1 - \beta_6)\left(z_t + \frac{1}{\rho}4\Delta y_t^*\right) + \vartheta_{6t} \tag{6}$$

$$z_t = \beta_7 z_{t-1} + \vartheta_{7t} \tag{7}$$

$$\widetilde{lev}_t = \beta_8\widetilde{lev}_{t-1} + \varphi_{81}\widetilde{dsr}_{t-1} + \varphi_{82}\widetilde{dsr}_{t-4} + \varphi_{83}(r_{t-1} - r_{t-1}^*) + \vartheta_{8t} \tag{8}$$

$$\widetilde{dsr}_t = \beta_9\widetilde{dsr}_{t-1} + \varphi_{91}\widetilde{lev}_{t-1} - \varphi_{92}\widetilde{lev}_{t-4} + \varphi_{93}\Delta i_{L,t-1} + \vartheta_{9t} \tag{9}$$

$$i_{L,t} = i_{L,t-1} - \varphi_{101}(i_{L,t-1} - i_{t-1} - \overline{spr}) + \vartheta_{10t}, \tag{10}$$

where $\vartheta_{it} \sim iid(0, \sigma_i^2)$ and the quarterly potential growth rate in (6) is annualized by multiplying it by 4.

Our extension relative to the baseline system is shown in grey. We first adjust the baseline system to allow for the possibility that the leverage gap contains information about the output gap.[7] If one thought of the corresponding equation, (3), as a typical reduced-form expenditure function, then the leverage could be regarded as

[7] An additional benefit of using an observable variable (the leverage gap) to relate monetary policy and real outcomes in (3) is that it helps to anchor potential output and the natural rate in the filter.

a proxy for the strength of the credit constraint. It enters with a negative sign because a positive leverage gap is associated with low asset prices, and hence low credit growth and correspondingly lower output.[8] The specification dovetails with previous work, in which the growth rates of credit and (residential) property prices were found to greatly improve the estimates of potential output (Borio, Disyatat, and Juselius 2014, 2017). The leverage gap can be seen as a more structured way of capturing the same information.

If the leverage gap turns out to be informative for the output gap (i.e., $\varphi_{32} > 0$), we also need an equation that characterizes its evolution. Since the leverage gap drives a wedge between actual and potential output, it is reasonable to relate it to deviations between the actual and natural real interest rates. Rates above the natural rate should decrease asset prices or output, which in turn increases the leverage gap. In addition, the debt service gap feeds negatively into asset price growth and, hence, boosts the leverage gap (Juselius and Drehmann 2015). Since the debt service effect is quite sluggish in the data, as a shortcut to modeling it with a more elaborate autoregressive structure we allow it to enter with both one- and four-quarter lags (equation (8)).

As Juselius and Drehmann (2015) show, there is a feedback loop between the leverage gap and the debt service gap running from new debt to the debt service and, in turn, from debt service to asset prices and thereby back to the leverage gap. And this loop underpins how the financial cycle interacts with the real economy. For instance, it can lead to persistent endogenous cycles that pose a significant challenge for monetary policy. To capture this two-way interaction, we allow the leverage gap to feed into the debt service gap in equation (9). Again, we include both the first- and fourth-quarter lags of this variable. The debt service gap is also directly related to changes in the average nominal interest rate on debt.

Clearly, the average nominal interest rate on debt can be closely influenced by monetary policy. We therefore postulate that the

[8]The effects from the leverage gap are unlikely to reflect more standard wealth effects. Juselius et al. (2016) show that the leverage gap has a stable positive effect on output growth via credit growth in a vector error-correction model (VECM). This effect does not change when the growth rates of various asset prices enter the model.

average lending rate is a weighted average of the past lending and policy rates with a constant, $\overline{spr}$, in the long run. This is expressed in equation (10).

3. Results

We now briefly describe the data used in estimating the system (3)–(10) and the key results.

3.1 Data

We use quarterly time series for the United States from 1985:Q1 to 2015:Q1. Output, y_t, is measured by real GDP. Our short-term interest rate, i_t, is the three-month federal funds rate, and we use the log change in the core personal consumption expenditure index as our inflation measure, π_t. To construct the leverage gap we also need a credit stock and a general asset price index. As a measure for the credit stock, cr_t, we use total credit from all sources to the private non-financial sector, obtained from the financial accounts. Our aggregate asset price index, $p_{A,t}$, is a weighted average of residential property prices, commercial property prices, and equity prices. The appendix provides details about its construction. The average lending rate on the stock of credit, $i_{L,t}$, is constructed by dividing private-sector interest payments, recorded in the national accounts, by the credit stock.[9]

3.2 Estimates

We estimate the parameters of the system (3)–(10) using a Bayesian approach with relatively weak priors. We assume that all β_j parameters follow the gamma distribution with mean 0.7 and standard deviation 0.2. To ensure that the output gap is strictly stationary, we restrict β_3 to lie in the interval between 0 and 0.95, while we allow the other autoregressive parameters to take any value in the

[9] We also account for financial intermediation services indirectly measured (FISIM)—an estimate of the value of financial intermediation services provided by financial institutions. When national accounts compilers construct the sectoral accounts, parts of interest payments are reclassified as payments for services and classified as output of the financial intermediation sector.

unit interval. For φ_{31}, φ_{32}, φ_5, φ_{85}, and φ_{101}, we assume the gamma distribution with mean 0.3 and standard deviation 0.2 and that they are positive. Since we split the cross-equational effects of the debt service and leverage gaps between lags 1 and 4, we set the prior means of φ_{81}, φ_{82}, φ_{91}, and φ_{92} to 0.15 with corresponding standard deviations of 0.1. There is no good guide for setting the prior of φ_{93}, but a straightforward regression suggests that its value is close to 10. Hence, we use this as prior mean with standard deviation 0.5. Rather than estimating the discount rate, ρ, we set it to 0.99, in line with the literature. We also fix the inflation target, π_t^*, at 2.

We use the inverse gamma distribution for the shock variances. We rely on historical variances of the first difference of HP-filtered output, inflation, the leverage gap, the debt service gap, and the average lending rate to calibrate the prior means of $\sigma_{4,t}^2$, $\sigma_{5,t}^2$, $\sigma_{8,t}^2$, $\sigma_{9,t}^2$, and $\sigma_{10,t}^2$ respectively, as well as on the variance of HP-filtered output growth (as a baseline for the natural rate) for the prior means of $\sigma_{6,t}^2$ and $\sigma_{7,t}^2$. We set the standard deviation to 0.5 for all of them. Furthermore, to ensure that the output gap captures conventional business-cycle frequencies, we fix its variance in relation to the variance of potential output.[10]

Before turning to the estimates of potential output and the natural rate of interest, it is worth dwelling on some of the key coefficient estimates (table 1).

The posterior estimates reveal that the leverage gap is an economically more important output gap driver than the standard real interest rate gap. The estimated coefficients (φ_{31} and φ_{32}) are broadly similar. But in the sample, the interest rate gap varies between –3.5 and +1.5 percentage points, whereas the leverage gap ranges from –20 to +27 percent, which translates into a –0.9 to +1.2 direct effect on the output gap. Moreover, the leverage gap is more persistent, leading to much higher long-run effects.

[10]In particular, we set the scaling parameter $\lambda = \sigma_{3,t}^2/\sigma_{4,t}^2$ in such a way that the ratio between the sample variance of the output gap and the acceleration of potential output is similar to that of the HP filter. This gives approximately the same frequency cutoff for the business cycle as in the HP case. The two filters coincide if $\beta_3 = \varphi_{31} = \varphi_{32} = 0$. See Borio, Disyatat, and Juselius (2014) for an in-depth discussion of these restrictions.

Table 1. Parameter Estimates for the Minimal Extension

Equation	Parameter	Loading On To	Prior	Prior Std.	Posterior	Posterior Std.
IS Curve, (3)	β_3	$y_{t-1} - y^*_{t-1}$	0.70	0.20	0.738	0.083
	φ_{31}	$r_t - r^*_t$	0.30	0.20	0.037	0.022
	φ_{32}	$\widetilde{lev}_t$	0.30	0.20	0.045	0.010
Phillips Curve, (5)	β_5	$\pi_{t-1} - \pi^*$	1.00	0.00	1.000	0.000
	φ_5	$y_t - y^*_t$	0.30	0.20	0.015	0.010
Natural Rate, (6)	β_6	r^*_{t-1}	0.70	0.20	0.616	0.069
z-factor, (7)	β_7	z_{t-1}	0.70	0.20	0.685	0.020
Leverage Gap, (8)	β_8	$\widetilde{lev}_{t-1}$	0.70	0.20	0.949	0.019
	φ_{81}	$\widetilde{dsr}_{t-1}$	0.15	0.10	0.004	0.001
	φ_{82}	$\widetilde{dsr}_{t-4}$	0.15	0.10	0.003	0.001
	φ_{83}	$r_t - r^*_t$	0.30	0.20	0.007	0.003
Debt Service, (9)	β_9	$\widetilde{dsr}_{t-1}$	0.70	0.20	0.960	0.011
	φ_{91}	$\widetilde{lev}_{t-1}$	0.15	0.10	0.020	0.006
	φ_{92}	$\widetilde{lev}_{t-4}$	0.15	0.10	0.051	0.008
	φ_{93}	$\Delta i_{L,t}$	10.00	0.50	7.442	0.310
Lending Rate, (10)	φ_{101}	$i_{L,t-1} - i_{t-1}$	0.30	0.20	0.093	0.012

Note: Results from estimating system (5)–(10) using a Kalman filter.

This suggests that the leverage gap is an important output gap driver and one of the main channels through which monetary policy can influence the real economy. But how effective is the real interest rate as a tool for controlling the gap? The answer is "not very." The coefficient estimate of the real interest rate gap in (8), φ_{83}, is about 0.007. Given the leverage gap range, this is tiny *on impact*. But given the leverage gap's high degree of persistence, the effect is eventually twenty times as large and implies that a 1 percentage point increase in the real rate above the natural rate increases the leverage gap by 0.13 percentage points. That said, the real interest rate effect on the output gap is much smaller than that reported in Laubach and Williams (2003). One reason might be that we use a shorter sample owing to the limited availability of commercial property prices.

The estimates of equation (8) also reveal that the lagged debt service gap terms are important drivers of the leverage gap. Given the range of the debt service gap, its economic effect is about ten times as large as that of the interest rate gap. The debt service gap, in turn, is also affected by the leverage gap. This gives rise to the potential for endogenous cycles (Juselius and Drehmann 2015). Moreover, via the effects on the lending rate (equation (10)), the nominal policy rate is directly linked to the debt service gap, implying a transmission channel from debt service to leverage and finally to output and inflation.

The estimates of the Phillips curve, (5), seem reasonable. The coefficient on lagged inflation is very close to unity. For simplicity, we impose this value strictly, in line with much of the literature. The coefficient on the output gap is rather small, but this is consistent with the literature pointing to a weak link between domestic slack and inflation. Finally, the estimates for the z_t factor suggest that this element is clearly stationary and relatively small.

Figure 3 compares our estimates of the output gap and the natural interest rate with those of Laubach-Williams (2015b).[11] Two points stand out with respect to the output gap (left-hand panel).

First, when estimated over the full sample, the two gaps move together from the mid-1990s. Recognizing the financial tailwinds, our output gap measure clearly indicates that the economy was running

[11]Note that these are two-sided, not real-time, estimates.

Figure 3. The Financial Cycle: Implications for the Natural Rate and Trend Output

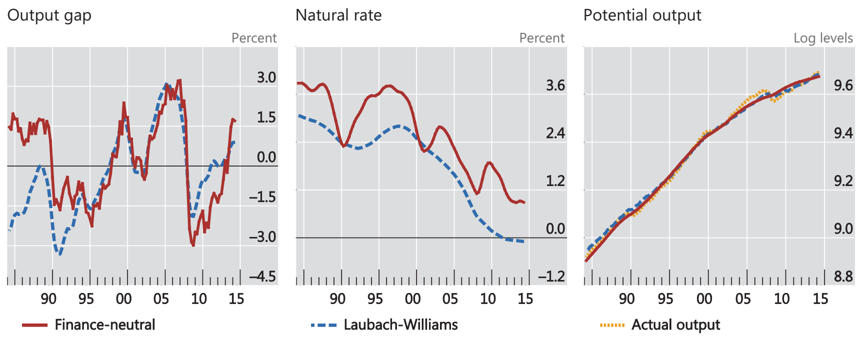

Sources: Laubach and Williams (2015b); national data; authors' calculations.

Notes: The finance-neutral variables are the result of estimating system (3)–(10). For the Laubach-Williams variables, we show the results of the two-sided filter using data until 2015:Q3 taken from Laubach and Williams (2015b).

above sustainable levels in the years leading up to the financial crisis. Conversely, output was below potential in the aftermath of the crisis owing to the substantial financial headwinds. The fact that output moved above potential towards the end of the sample reflects the significant support that financial factors provided to the U.S. economy during that phase, with leverage and debt service below their long-run levels.

Second, in contrast to the Laubach-Williams output gap, which is persistently negative during most of the 1980s and 1990s, ours is positive ahead of the recession in the early 1990s and only negative afterwards. This is because a financial boom was under way at the time, qualitatively similar to the one that preceded the more recent crisis but smaller. Indeed, some banks faced serious strains in the early 1990s, and the expression "financial headwinds" was quite common (Greenspan 2004).

A variance decomposition also provides a different perspective on the drivers of the output gap than what the standard literature would suggest. In particular, inflation contributes next to nothing to the variance in the output gap in our specification (left-hand panel, figure 4). The leverage gap is the main contributor, followed by the real interest rate gap.

Figure 4. The Financial Cycle Helps Explain the Variation in the Output Gap and the Natural Rate (variance decomposition of the output gap and the natural rate, in percent)

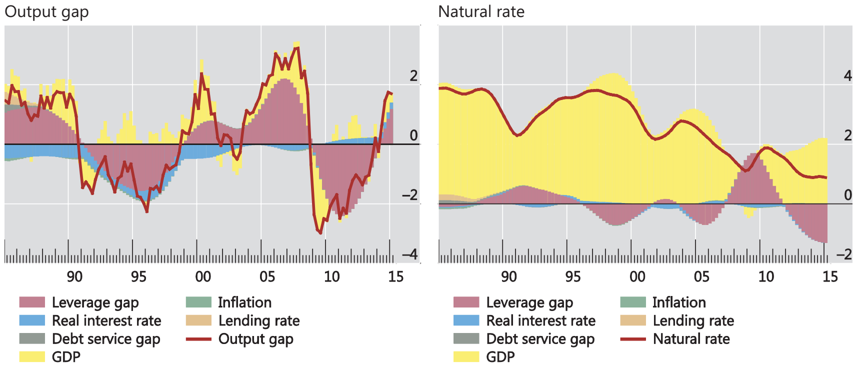

Sources: National data; authors' calculations.

For present purposes, the focus is on the natural interest rate (middle panel, figure 3). In particular, our estimate shows a decline from 4 percent to 0.9 percent over the last thirty years. This is in line with the downward trend in potential output growth (right-hand panel, figure 3). Starting in 2000, our natural rate estimate has clearly declined by less than its Laubach-Williams counterpart. Interestingly, actual interest rates (see figure 1) have remained below the estimated natural rate for almost the whole period under study. Sharp interest cuts in response to financial strains in the early 1990s, early 2000s, and 2008 were not taken back in the ensuing normalization phase, suggesting substantial policy asymmetry with respect to the financial cycle. Finally, given the tight specification, GDP growth accounts for most of the variance in the evolution of the natural rate (right-hand panel, figure 4).

By contrast, the natural rate estimated by Laubach-Williams is consistently below our estimates and is currently negative (middle panel, figure 3). This reflects the emphasis the framework places on inflation in pinning down both the natural rate and potential output, through the IS curve and the Phillips curve, respectively. In the early part of the sample, the downward trend in inflation leads to

Figure 5. A Comparison of Various Natural Rate Estimates (in percent)

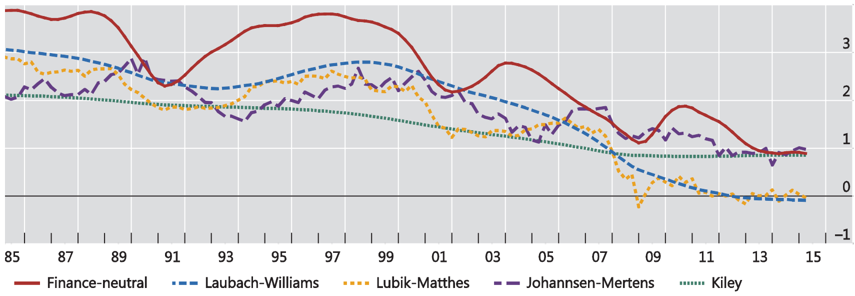

Sources: Laubach and Williams (2015b); Lubik and Matthes (2015); Kiley (2015); Johannsen and Mertens (2016); authors' calculations.

a persistently negative output gap and thus to an estimate of the natural rate that is generally below the real interest rate. Since the mid-90s, the Laubach-Williams estimates of the output gap fluctuate around zero. And given the IS curve, which depends only on the interest rate gap, the estimated natural rate must be below the observed real rate whenever the output gap is negative. In contrast, our estimates are less sensitive to these factors, as the Phillips curve and the interest rate gap in the IS curve do not play such prominent roles.

For broader comparison, figure 5 plots our estimates of the finance-neutral natural rate alongside those constructed in a number of recent papers. For instance, Lubik and Matthes (2015) derive an estimate from a time-varying parameter vector autoregressive model for real GDP, inflation, and the real interest rate. Despite the sharp difference in method, they obtain estimates that are remarkably similar to those of Laubach and Williams. By contrast, using a framework that is similar in spirit to ours, Kiley (2015) finds that estimates of the natural rate become much higher once one augments the IS curve with additional financial variables, such as a corporate bond spread and credit growth. While the resulting estimates fluctuate less than ours and are generally lower, the estimated levels are very close in the recent period. A similar finding also appears in Johannsen and Mertens (2016), who estimate the expected long-run value for the real federal funds rate from an empirical model of the shadow rate that conditions on a long-term interest rate and the Congressional Budget Office unemployment gap.

Finally, it should be noted that our filtering system is quite robust in real time, confirming previous work on output gaps using financial information (Borio, Disyatat, and Juselius 2014, 2017). Coefficients hardly change, so that we obtain very similar output gap and natural rate estimates even after observing more data (Juselius et al. 2016). This is a very desirable property from a policymaking perspective.

4. The Financial Cycle and Monetary Policy

The above filtering system suggests that much of the cyclical movements in output can be attributed to financial factors. Moreover, the strong feedback dynamics between the leverage gap and the debt service gap suggest that the current state of the business cycle alone may be too crude a guide for policy. Recognizing this should lead to a different design of stabilization policies, not least that of monetary policy. In this section we take a preliminary step in exploring this. To do so, we propose a modification to a standard monetary policy rule, by including a systematic response to financial developments. We then evaluate our policy rule in an illustrative counterfactual experiment.

4.1 An Alternative Monetary Policy Rule

As an alternative policy rule, we start from the popular Taylor rule and change it in two respects. First, consistent with our analysis, we allow the natural rate to change over time, i.e., the intercept in the rule is no longer a constant. This is quite standard.[12] Second, we augment it with a financial-cycle indicator. The simplest way of doing this is to add the debt service gap. As the filter shows, the

[12]Carlstrom and Fuerst (2016) analyze this formally, in particular within a standard New Keynesian model with $r_t^* = r^* + \Delta prod_t$, where $\Delta prod$ is expected productivity growth. They then assess Taylor rules of the form $i_t = (r^* + \alpha(r_t^* - r^*) + \pi^* + 1.5(\pi_t - \pi^*) + 0.1(y_t - y_t^*))$.

debt service gap closely influences the leverage gap and hence the output gap (see equations (8) and (3)). And in comparison with the leverage gap, the debt service gap is strongly influenced by policy through lending rates (equations (9) and (10)). Thus, including the debt service gap in the policy rule is one way to increase traction over the financial cycle. Hence, we analyze the following rule:

$$\begin{aligned} i_t = \rho i_{t-1} + (1-\rho)(r^*_{t-1} + \pi^* + 1.5(\pi_{t-1} - \pi^*) \\ + 0.5(y_{t-1} - y^*_{t-1}) - \lambda \, \widetilde{dsr}_{t-1}). \end{aligned} \tag{11}$$

Here, we set $\rho = 0.85$ and $\lambda = 0.75$. In Juselius et al. (2016), we also explored other values for λ, including $\lambda = 0$, so that our rule collapses into the standard Taylor rule with a time-varying intercept that equals the natural rate.

4.2 A Counterfactual Experiment

To assess the potential benefits from monetary policy taking into account the financial cycle, we conduct a counterfactual exercise in which we embed the new policy rule in the economy as estimated by the filter. We essentially ask what the evolution of the economy would have been from a given point in time if (i) policymakers followed the policy rule in (11), (ii) agents' behavior remained invariant, so that the reduced-form structure of the economy estimated by the filter did not change, and (iii) the economy was hit by the same historical shocks, including for instance the financial crisis, but excluding shocks to monetary policy. Put differently, differences between the counterfactual and historical outcomes are solely due to the systematic policy interest rate path.

The assumption that agents' behavior does not change as we change the systematic policy response is obviously at odds with the Lucas critique. This concern should not be understated. That said, we draw some comfort from past studies that have found that the Lucas critique may be of limited relevance in practice.[13] For instance,

[13]Relaxing some of the strong assumptions that underpin mainstream monetary policy models could weaken the force of Lucas's argument even theoretically. For example, incorporating features such as rule-of-thumb agents, model uncertainty, ambiguity, incomplete information, multiple equilibria, or constrained agents can have this effect.

a common finding is that the parameters of empirical VARs are remarkably stable despite changes in estimated policy equations in the sample (e.g., Favero and Hendry 1992, Leeper and Zha 2003, Rudebusch 2005).[14] In addition, in Juselius et al. (2016) we implement an analogous counterfactual analysis with a slightly different approach and obtain very similar, in fact stronger, results.

Nevertheless, given that our counterfactual alters the monetary policy rule rather than just its coefficients, the Lucas critique may have more force in this context. Even then, at least two aspects are worth highlighting.

First, the Federal Reserve has seemingly reacted to debt service burdens in the past. For instance, it explicitly took debt service burdens into account when setting policy under Greenspan (Greenspan 1993). In addition, in Juselius et al. (2016) we find some evidence that it has, directly or indirectly, reacted to a high debt service burden also more recently. If so, our policy experiment would involve more a change in the intensity of the policy response than a fundamental change in the reaction function's shape.

Second, in some respects, an explicit acknowledgement of the Lucas critique might even strengthen our results. As explained below, monetary policy potentially has a large impact on output dynamics through its influence on the financial cycle. Were market participants to internalize the systematic response to financial developments, the policy's effectiveness in dampening the financial cycle could arguably be greater.

Despite these arguments, we fully acknowledge the shortcomings of our counterfactual exercise. On balance, given the potential changes in behavior, our results can at best be seen as giving a rough indication of the benefits of a policy shift—a preliminary step that will need to be corroborated by further research and different approaches.

[14] Linde (2001) and Lubik and Surico (2010) argue that these findings are due to the weak power of the stability tests. They find that changes in policy led to corresponding changes in the VAR parameters. But even if such changes can be detected statistically, they do not seem to be very large economically.

To implement the counterfactual experiment, we follow an iterative procedure, starting from 2003:Q1 as t_0:[15]

(i) Derive the natural rate $r^*_{t_0-1}$ and the output gap $(y_{t_0-1} - y^*_{t_0-1})$ using the estimated filter.
(ii) Set policy rate for t_0 as $i_{t_0} = 0.85\rho + 0.15(r^*_{t_0-1} + \pi^* + 1.5(\pi_{t_0-1} - \pi^*) + 0.5(y_{t_0-1} - y^*_{t_0-1}) - 0.75\widetilde{dsr}_{t_0-1})$ if this leads to $i_{t_0} > 0$ or set $i_{t_0} = 0$ otherwise.
(iii) Use the estimated filter and generate predictions of all variables in the system for time t_0 conditional on the new policy rate and the retained errors ε_{t_0}.
(iv) Redo steps (i)–(iii) for $t_0 + 1, t_0 + 2 \ldots$ until the end of the sample.

The various caveats notwithstanding, the counterfactual exercise shows that the alternative policy rule potentially yields considerable output gains compared with actual history with little change in inflation. By the end of the simulation period the cumulative output gain is over 15 percent, or around 1.25 percent per year (top left-hand panel, figure 6). As both the debt service gap and the leverage gap are initially negative—the latter strongly so—the policy rule calls for leaning against the financial boom by raising rates (second row, figure 6). Initially, the negative leverage gap and higher rates drive up debt service, which in turn helps to stabilize the leverage gap. But higher policy rates also weigh on economic activity. In fact, output in the counterfactual policy prior to 2008:Q4 is on average 1 percent lower than actually experienced. These output losses, however, are more than offset by the gains during the Great Recession, which still occurs also because the Lehman shock is included in the counterfactual. Moreover, lower financial imbalances in the form of lower debt service and leverage gaps mean that output after 2008 is on average 3 percent higher in the counterfactual than historically (third row, figure 6). And the leeway created by following the policy rule pre-crisis also means that interest rates are on average higher. As such, the central bank retains greater room for policy maneuver.

[15]In Juselius et al. (2016), we also explore a counterfactual which starts in 1996. We find that an earlier implementation succeeds in containing financial imbalances much better, thereby leading to higher gains.

Figure 6. Leaning Against the Financial Cycle Improves Economic Performance

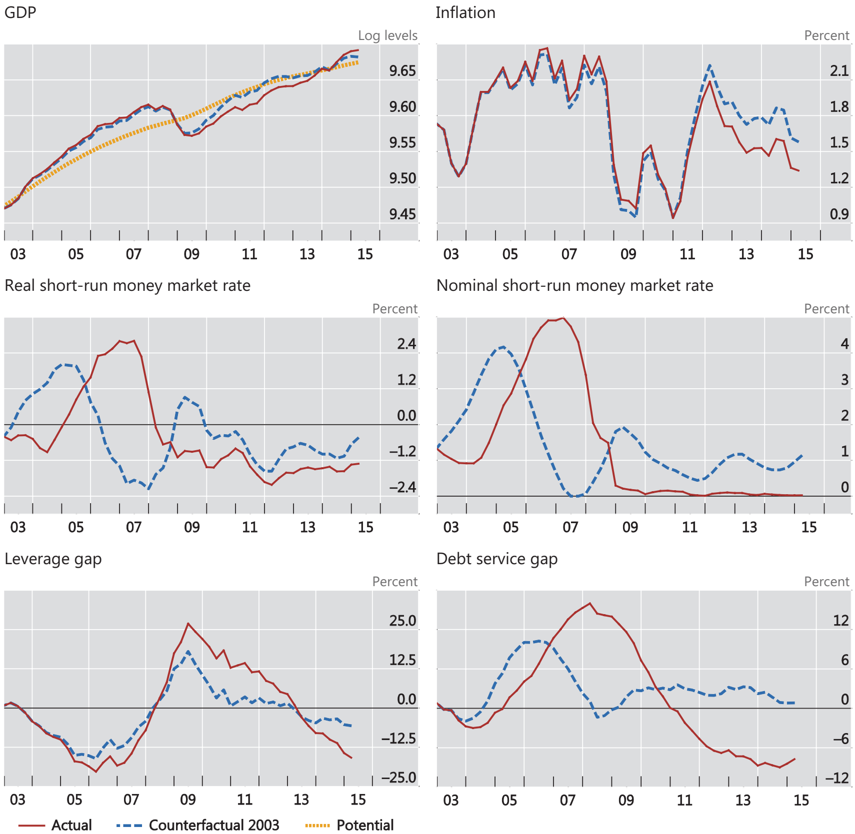

Sources: National data; authors' calculations.

Notes: In the counterfactual experiment, we set policy in line with an augmented Taylor rule that takes account of the finance-neutral natural rate, the finance-neutral output gap, and the debt service gap in line with equation (11). Results are based on the filter (3)–(10). We retain the historical errors to derive the evolution of the variables in the counterfactual. The counterfactual policy starts in 2003:Q1.

These output gains come with little change in overall inflation performance. This is not too surprising given the low traction that the output gap has on inflation in the filter. In fact, the output gap is, on average, smaller in the counterfactual. In particular, after

2010, output is close to potential and inflation ends up being a bit *higher* and closer to target, not lower.

It should be noted that retaining the Lehman crisis residuals stacks the deck against us. Presumably, had the authorities succeeded in restraining the boom in the first place, the shock might have been smaller, and might not even have materialized. This should be considered when evaluating the findings.

Finally, by making minimal changes to the standard filtering framework, we have deliberately restricted the extent to which the financial cycle, and hence monetary policy, can affect longer-run output trajectories. In particular, potential output and the natural rate are fully exogenous in the filter and do not change with changes in monetary policy. Moreover, since we assume that the estimated shocks from the filter also occur in the counterfactual experiment, both of these variables will take exactly the same paths as in the full sample. In the more elaborate, flexible, and in our view more realistic, setup detailed in Juselius et al. (2016), where economic trajectories are modeled through a VECM, not only are output gains under the alternative policy rule more persistent, but the natural rate estimates in the counterfactual are also higher. Thus, the results presented in this section in some sense can be viewed as a conservative illustration of the extent to which the monetary policy–finance nexus can influence macroeconomic outcomes.

4.3 Policy Considerations

Our analysis points to a number of possible shortcomings of the typical *empirical* framework employed to consider the benefits and costs of a "leaning against the wind" policy intended to reduce financial instability risks. In that framework, policy is calibrated to reduce the probability of a crisis by deviating temporarily from its usual systematic response to influence a variable, typically credit growth, found empirically to have good leading-indicator properties for banking crises (e.g., Svensson 2014, 2016; IMF 2015).

First, the typical empirical framework understates the costs of financial imbalances to the extent that it ignores persistent, and possibly permanent, effects on the *level* of output. Our analysis above shows that the output effects even in the simple system can be quite persistent. In the more flexible framework of Juselius et al. (2016), the impact of the financial cycle on output lasts much longer, which is more consistent with the evidence that crises may have permanent effects on the level of output (Cerra and Saxena 2008, BCBS 2010, Ball 2014). Moreover, these costs may arise even if a full-blown crisis does not occur. Hence, the costs of neglecting the financial cycle are likely to be an order of magnitude higher.

Second, the typical empirical framework underestimates the contribution of monetary policy to the imbalances. This is because it focuses on its *marginal* effect on the variables of interest, typically credit growth, but ignores its *cumulative* impact, notably on the credit-to-GDP ratio, and hence, through them, on the economy's path.[16] In addition, whenever the debt service ratio is ignored, the relevance of monetary policy is understated further, given its first-order effect on interest payments.

Third, for much the same reasons, the typical empirical framework can be misleading. Thinking of a "leaning against the wind" policy as one that involves temporary deviations from an otherwise standard rule is not that helpful. What matters is the *systematic* policy followed along the whole financial cycle, i.e., avoiding straying persistently too far away from financial equilibrium, with large buildups in the two financial gaps. Following a "business as usual" policy most of the time, combined with occasional leaning only once the signs of financial imbalances become obvious, would result in doing too little too late. At worst, the central bank could simply be seen as precipitating the very recession it wishes to prevent. Selective attention is not the answer. A "through-the-cycle" policy is called for.[17]

[16] One could raise a similar objection to the leading indicator of the crisis itself—credit growth. It is not credit growth per se that provides a good signal but the *cumulative* growth over and above certain thresholds alongside other developments, such as abnormal increases in asset prices (Borio and Drehmann 2009) or the behavior of the debt service burden (Drehmann and Juselius 2013). For example, periods of rapid credit growth early on in the cycle are unlikely to signal impending crises. All of this introduces "noise" in the indicator's predictive content. This, in turn, will inevitably reduce the benefits of a leaning policy.

[17] See also Filardo and Rungcharoenkitkul (2016), who reach broadly similar conclusions based on a model that highlights the importance of the systematic policy rule when the financial cycle is endogenous.

At the same time, we should not take the suggested rule too literally. For instance, looking more closely at the behavior of the underlying variables, the policy would have called for starting to ease around the peak in property prices and when credit expansion was still rather strong. This is because at that point the debt service gap switches sign, moving above its long-term average. It may well take a brave central bank, with great confidence in the underlying relationships, to stop tightening under those conditions. This puts a premium on the use of complementary tools, such as macroprudential measures, in the later stages of financial booms.

The analysis highlights the risks of policies that are asymmetrical in relation to the financial cycle. This is, in effect, what a policy focused primarily on inflation and short-term output fluctuations can produce, as the U.S. example here illustrates—simply one among many. Such a policy does little, if anything, to restrain the upswing but reacts strongly and persistently to the downswing. In the case in question, it translates most conspicuously into an asymmetry in the evolution of the debt service gap, which was positive on average over the sample from 2000, ultimately resulting in a lower output path.

The risks involved are apparent. Over time, such an asymmetrical policy can impart a downward bias to interest rates as the buildup of debt over successive boom-bust cycles leads to depressed economic activity, making it increasingly hard to raise interest rates—a kind of "debt trap" (Borio and Disyatat 2014, Borio 2016). That is, both the leverage and debt service gaps end up being significantly above their long-run equilibrium levels and the growth impetus from already low interest rates is limited. At this point, the economy is overindebted and overleveraged, making it difficult to raise rates without damaging it.

These considerations highlight the possibility that, in more ways than one, over long horizons low interest rates may become, to some extent, self-validating. Low rates may beget lower rates as monetary policy contributes to financial booms and busts. And to the extent that these forces exert a temporary, if potentially persistent, impact on potential output growth, the natural rate may also be affected. Either way, policy rates would not be just passively reflecting some deep exogenous forces; they would also be helping to shape the economic environment policymakers take as given ("exogenous") when tomorrow becomes today. Path dependence is key. Unless financial factors are taken more systematically into account, and a sufficiently long horizon is adopted, policy steps that appear reasonable when taken in isolation may take policy astray when considered as a sequence. Central banks need to carefully weigh the medium- to long-term side effects of policy working through the financial cycle against the benefits of short-term stimulus.

5. Conclusion

The critical role that financial developments play in economic fluctuations has long been recognized. Yet the prevailing analysis of the business cycle, and of its relationship to interest rates, does not exploit these interlinkages much. The extraction of trends and long-run equilibrium variables, such as potential output and the natural interest rate, need to go beyond the standard full employment-inflation paradigm. Surely, equilibrium outcomes should also be sustainable. If the ebb and flow of the financial cycle coincides with damaging economic booms and busts, then assessments of the sustainability of a given path for output or interest rates need to take financial developments into account. Financial and macroeconomic stability are essentially two sides of the same coin.

In contrast to the prevailing view, we argue that an exogenous decline in the natural real interest rate provides an incomplete explanation of the observed trend reduction in real interest rates and of their persistence at ultra-low levels today. Based on U.S. data, we find that if one considers the influence of the financial cycle, the estimated natural interest rate is generally higher and, on balance, has declined by less since the 2000s. Moreover, policy rates have been persistently below this estimate. We also find that monetary policy has a first-order effect on the financial cycle and that financial busts can have very persistent effects on output. Together, these findings indicate that part of the observed decline in market interest rates reflects the interaction between monetary policy and the financial cycle. They suggest that policy has leaned aggressively and persistently against financial busts, but has failed to lean sufficiently promptly and deliberately against financial booms. The resulting asymmetry appears to have contributed to a downward bias in interest rates.

Accordingly, an illustrative counterfactual experiment suggests that a policy rule that systematically takes into account financial developments helps to dampen the financial cycle, leading to higher output with little change in inflation. Such a policy can also result in a smaller decline in the estimated natural rate. Because of well-known econometric limitations, this part of the analysis should be interpreted with great caution. At a minimum, though, it indicates that it is inappropriate to think of a financial-stability-oriented monetary policy as one that simply leans against signs of the buildup of financial imbalances only when they become evident. Such a "selective attention" strategy could easily result in doing too little too late and would likely backfire. Rather, the right policy would need to take financial considerations *systematically* into account, never straying too far away for too long from some notion of "financial equilibrium." We conjecture that this conclusion, and the merits of such a policy more generally, will withstand further scrutiny.

Clearly, our analysis is just one small further step in the development of a monetary policy framework that takes financial stability considerations, broadly defined, into account. For one, rather than being based on a fully fledged "structural" model, it hinges on some key statistical relationships found in the data. While these could in principle be derived from more fundamental behavioral relationships and embedded in a system better suited for counterfactual policy analysis, we leave this for future work. Similarly, the econometric findings would be more convincing if they were shown to hold both across countries and monetary policy regimes. This would go a considerable way in addressing also the Lucas critique. We leave this, too, for future work.

Despite the limitations of our analysis, we hope to have shown that it is possible to make further progress in making a financial-stability-oriented monetary policy framework more operational. And as argued elsewhere (Bank for International Settlements 2014, 2015), recognizing this could help integrate monetary policy into a more holistic and balanced macrofinancial stability framework that would include also other policies, notably prudential and fiscal policies. This would be a more effective way of promoting *lasting* monetary, financial, and macroeconomic stability.

Appendix. The Leverage and Debt Service Gaps

In this appendix, we discuss the construction and estimation of the leverage and debt service gaps.

To accurately measure the debt service gap, we need the average interest rate on the stock of debt. This variable can be obtained from the national accounts as gross interest payments plus financial intermediation services indirectly measured divided by the stock of credit (Drehmann et al. 2015).

For the leverage gap we need an aggregate asset price index. This is generally not available and has to be constructed from the price indexes of various sub-asset classes. Given stark differences between the types of collateral used by households (HH) and firms (NFC), as well as changes in the sectoral composition over the sample, we first estimate separate asset price indexes for each sector separately from long-run empirical relationships of the form

$$\widetilde{lev}_t^s = c_t^s -{}_t^s -\alpha_1^s p_{res,t}^r - \alpha_2^s p_{com,t}^r - (1 - \alpha_1^s - \alpha_2^s) p_{eq,t}^r - \mu_{lev}^s,$$

where $p_{res,t}$ denotes residential property prices, $p_{com,t}$ represents commercial property prices, and $p_{eq,t}$ stands for equity prices. These asset prices are obtained from the BIS databases. Superscript $s = H, NFC$ is used to denote the type of agent and superscript r is used to denote real values (e.g., $p_{res,t}^r = p_{res,t} - p_t$). We then construct the aggregate asset price as a weighted average

$$p_{A,t}^r = \omega_t p_{A,t}^{r,HH} + (1 - \omega_t) p_{A,t}^{r,NFC},$$

where ω_t is the share of private non-financial credit in the household sector.

We estimate the long-run parameters α_j^s for each sector from separate VAR models in error-correction form. Using Johansen's likelihood ratio test, we find that the co-integration rank is one in both cases. Moreover, formal exclusion tests suggest that commercial property prices and equity prices can be excluded for the households, whereas residential property prices can be excluded for the NFCs. Doing so gives us the estimates, $p_{A,t}^{r,HH} = p_{res,t}^r$ and $p_{A,t}^{r,NFC} = 0.72 p_{com,t}^r + 0.28 p_{eq,t}^r$. By a similar approach, we obtain the estimate $\beta_{dsr} = 5.54$ in (2). The details of these estimations can be found in Juselius et al. (2016).

Juselius et al. (2016) investigate the robustness of the estimates with respect to different samples. Despite their relatively high persistence, the two gaps can be accurately estimated from any sample that uses data up to at least 2003. The difference between the full-sample leverage gap and the gaps estimated over different subsamples is at most 2.5 percentage points, except in the mid-1980s where the difference is slightly higher. The stability of the debt service gap is even greater: the corresponding difference is at most 1.6 percentage points. Crucially, this implies that the gaps do not depend on the Great Recession and the financial boom that preceded it.

In Juselius et al. (2016), we also compare our estimated gaps with purely data-based measures as given by the credit-to-assets ratio for the financial accounts and debt service ratios calculated using a Federal Reserve Board methodology (Dynan, Johnson, and Pence 2003). The correspondence between these measures is close, suggesting that they broadly capture the same information.[18]

References

Adrian, T., and F. Duarte. 2016. "Financial Vulnerability and Monetary Policy." Staff Report No. 804, Federal Reserve Bank of New York.

Ajello, A., T. Laubach, D. López-Salido, and T. Nakata. 2015. "Financial Stability and Optimal Interest-Rate Policy." Paper presented at "The New Normal for Monetary Policy" conference, held by the Federal Reserve Bank of San Francisco, March 27.

Arseneau, D., and M. Kiley. 2014. "The Role of Financial Imbalances in Assessing the State of the Economy." FEDS Notes No. 2014-04-18, Board of Governors of the Federal Reserve System.

Ball, L. 2014. "Long-Term Damage from the Great Recession in OECD Countries." *European Journal of Economics and Economic Policies* 11 (2): 149–60.

[18]See also annex 1 in Juselius and Drehmann (2015). To be precise, while the debt service ratio is stationary, the debt-to-assets ratio of households and non-financial corporates from the flow of funds exhibits a slight deterministic trend, likely due to the fact that asset prices in the financial accounts are not fully marked to market. Deviations from this linear trend are closely correlated with the deviations from the long-term co-integrating relationship.

Ball, L., and S. Mazumder. 2011. "Inflation Dynamics and the Great Recession." NBER Working Paper No. 17044.

Bank for International Settlements. 2014. *84th BIS Annual Report* (June).

———. 2015. *85th BIS Annual Report* (June).

Basel Committee on Banking Supervision. 2010. "An Assessment of the Long-Term Economic Impact of Stronger Capital and Liquidity Requirements" (July).

Bean, C., C. Broda, T. Ito, and R. Kroszner. 2015. "Low for Long? Causes and Consequences of Persistently Low Interest Rates." Geneva Reports on the World Economy No. 17, ICMB and CEPR.

Bernanke, B. 2015. "Why Are Interest Rates So Low? Part 3: The Global Savings Glut." Brookings Blog Entry, April 1. Available at http://www.brookings.edu/blogs/ben-bernanke/posts/2015/04/01-why-interest-rates-low-global-savings-glut.

Bernanke, B. S., M. Gertler, and S. Gilchrist. 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." In *Handbook of Macroeconomics*, Vol. 1C, ed. J. B. Taylor and M. Woodford, 1341–93 (chapter 21). Elsevier Science.

Blanchard, O., E. Cerutti, and L. Summers. 2015. "Inflation and Activity—Two Explorations and their Monetary Policy Implications." NBER Working Paper No. 21726.

Borio, C. 2016. "Revisiting Three Intellectual Pillars of Monetary Policy." *Cato Journal* 36 (2): 213–38. Also available at http://www.bis.org/speeches/sp151112.htm.

Borio, C., and P. Disyatat. 2014. "Low Interest Rates and Secular Stagnation: Is Debt a Missing Link?" VoxEU, June 25.

Borio, C., P. Disyatat, and M. Juselius. 2014. "A Parsimonious Approach to Incorporating Economic Information in Measures of Potential Output." BIS Working Paper No. 442 (February).

———. 2017. "Rethinking Potential Output: Embedding Information about the Financial Cycle." *Oxford Economic Papers* 69 (3): 655–77.

Borio, C., and Drehmann. 2009. "Assessing the Risk of Banking Crises—Revisited." *BIS Quarterly Review* (March): 29–46.

Borio, C., and A. Filardo. 2007. "Globalisation and Inflation: New Cross-Country Evidence on the Global Determinants of Domestic Inflation." BIS Working Paper No. 227.

Borio, C., and P. Lowe. 2002. "Asset Prices, Financial and Monetary Stability: Exploring the Nexus." BIS Working Paper No. 114 (July).

Broadbent, B. 2014. "Monetary Policy, Asset Prices and Distribution." Speech given at the Society of Business Economists Annual Conference, London, October 23.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas. 2008. "An Equilibrium Model of 'Global Imbalances' and Low Interest Rates." NBER Working Paper No. 11996.

Carlstrom, C. T., and T. Fuerst. 2016. "The Natural Rate of Interest in Taylor Rules." Economic Commentary No. 2016-01, Federal Reserve Bank of Cleveland.

Cerra, V., and S. Saxena. 2008. "Growth Dynamics: The Myth of Economic Recovery." *American Economic Review* 98 (1): 439–57.

Claessens, S., M. Kose, and M. Terrones. 2009. "What Happens during Recessions, Crunches and Busts?" *Economic Policy* 24 (60): 655–700.

Cúrdia, V., and M. Woodford. 2016. "Credit Frictions and Optimal Monetary Policy." *Journal of Monetary Economics* 84: 30–65.

Del Negro, M., D. Giannone, M. Giannoni, and A. Tambalotti. 2017. "Safety, Liquidity, and the Natural Rate of Interest." *Brookings Papers on Economic Activity* (Spring).

Del Negro, M., M. Giannoni, and F. Schorfheide. 2015. "Inflation in the Great Recession and New Keynesian Models." *American Economic Journal: Macroeconomics* 7 (1): 168–96.

Dotsey, M., S. Fujita, and T. Stark. 2015. "Do Phillips Curves Conditionally Help to Forecast Inflation?" Working Paper No. 15-16, Federal Reserve Bank of Philadelphia.

Drehmann, M., A. Ilnes, M. Juselius, and M. Santos. 2015. "How Much Income is Used for Debt Payments? A New Database for Debt Service Ratios." *BIS Quarterly Review* (September): 89–103.

Drehmann, M., and M. Juselius. 2013. "Evaluating Early Warning Indicators of Banking Crises: Satisfying Policy Requirements." BIS Working Paper No. 421 (August).

Dynan, K., K. Johnson, and K. Pence. 2003. "Recent Changes to a Measure of U.S. Household Debt Service." *Federal Reserve Bulletin* 89 (10): 417–26.

Faust, J., and E. Leeper. 2015. "The Myth of Normal: The Bumpy Story of Inflation and Monetary Policy." Paper prepared for the 39th Annual Economic Policy Symposium, sponsored by the Federal Reserve Bank of Kansas City, held in Jackson Hole, Wyoming, August 27–29.

Favero, C., and D. Hendry. 1992. "Testing the Lucas Critique: A Review." *Econometric Reviews* 11 (3): 265–306.

Filardo, A., and P. Rungcharoenkitkul. 2016. "Quantitative Case for Leaning against the Wind." BIS Working Paper No. 594.

Greenspan, A. 1993. "Semiannual Monetary Policy Report to the Congress." Testimony before the Committee on Banking, Housing and Urban Affairs, U.S. Senate, February 19.

———. 2004. "Risk and Uncertainty in Monetary Policy." Remarks at the meetings of the American Economic Association, San Diego, California, January 3.

International Monetary Fund. 2014. *World Economic Outlook: Recovery Strengthens, Remains Uneven* (April).

———. 2015. "Monetary Policy and Financial Stability." IMF Staff Report.

Johannsen, B., and E. Mertens. 2016. "A Time Series Model of Interest Rates with the Effective Lower Bound." Finance and Economics Discussion Series No. 2016-033, Board of Governors of the Federal Reserve System.

Jordà, Ò., M. Schularick, and A. M. Taylor. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking* 45 (s2): 3–28.

Juselius, M., C. Borio, P. Disyatat, and M. Drehmann. 2016. "Monetary Policy, the Financial Cycle and Ultra-Low Interest Rates." BIS Working Paper No. 569.

Juselius, M., and M. Drehmann. 2015. "Leverage Dynamics and the Real Burden of Debt." BIS Working Paper No. 501.

Kiley, M. 2015. "What Can the Data Tell Us about the Equilibrium Real Interest Rate?" Finance and Economics Discussion Series No. 2015-077, Board of Governors of the Federal Reserve System (August).

Laubach, T., and J. C. Williams. 2003. "Measuring the Natural Rate of Interest." *Review of Economics and Statistics* 85 (4): 1063–70.

———. 2015a. "Measuring the Natural Rate of Interest Redux." Working Paper No. 16, Federal Reserve Bank of San Francisco.

———. 2015b. "Real-time Estimates of Laubach-Williams Model." Available at www.frbsf.org/economic-research/economists/john-williams/Laubach_Williams_real_time_estimates_2005_2015.xlsx.

Leeper, E. M., and T. Zha. 2003. "Modest Policy Interventions." *Journal of Monetary Economics* 50 (8): 1673–1700.

Linde, J. 2001. "Testing for the Lucas Critique: A Quantitative Investigation." *American Economic Review* 91 (4): 986–1005.

Lubik, T., and C. Matthes. 2015. "Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches." Economic Brief No. 15-10 (October), Federal Reserve Bank of Richmond.

Lubik, T., and P. Surico. 2010. "The Lucas Critique and the Stability of Empirical Models." *Journal of Applied Econometrics* 25 (1): 177–94.

Mian, A., A. Sufi, and E. Verner. 2016. "Household Debt and Business Cycles Worldwide." NBER Working Paper No. 21581.

Obstfeld, M., and L. Tesar. 2015. "The Decline in Long-Term Interest Rates." Report, Council of Economic Advisers. Available at https://www.whitehouse.gov/blog/2015/07/14/decline-long-term-interest-rates.

Pain, N., I. Koske, and M. Sollie. 2008. "Globalisation and OECD Consumer Price Inflation." *OECD Economic Studies* 44 (1): 1–32.

Rachel, L., and T. Smith. 2015. "Secular Drivers of the Global Real Interest Rate." Staff Working Paper No. 571, Bank of England.

Reinhart, C., and K. Rogoff. 2009. *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press.

Rudebusch, G. 2005. "Assessing the Lucas Critique in Monetary Policy Models." *Journal of Money, Credit and Banking* 37 (2): 245–72.

Schularick, M., and A. Taylor. 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008." *American Economic Review* 102 (2): 1029–61.

Summers, L. 2014. "U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound." *Business Economics* 49 (2): 65–73.

Svensson, L. 2014. "Inflation Targeting and 'Leaning against the Wind'." *International Journal of Central Banking* 10 (2): 103–14.

———. 2016. "Cost-Benefit Analysis of Leaning against the Wind: Are Costs Larger Also with Less Effective Macroprudential Policy?" IMF Working Paper No. 16/3.

Taylor, J. 1993. "Discretion versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy* 39: 195–214.

Wicksell, K. 1898. *Geldzins und Güterpreise. Eine Untersuchung über die den Tauschwert des Geldes bestimmenden Ursachen.* Jena: Gustav Fischer (tr., 1936: *Interest and Prices. A Study of the Causes Regulating the Value of Money*, London: Macmillan).

Woodford, M. 2003. *Interest and Prices.* Princeton, NJ: Princeton University Press.

Discussion of “Monetary Policy, the Financial Cycle, and Ultra-Low Interest Rates”*

Marc P. Giannoni
Federal Reserve Bank of New York

1. Introduction

Several recent papers have documented a trend decline in real interest rates at least since the late 1990s in the United States and other countries (see, e.g., Holston, Laubach, and Williams 2017, among others). While monetary policy has slashed the federal funds rate in the wake of the Great Recession and kept it near zero until the end of 2015, researchers have argued that the low rates reflect in large part a depressed real “natural rate of interest,” r^*, i.e., an interest rate that would prevail if output were equal to its potential. Such an interest rate, which abstracts from monetary policy, captures the effects of real forces that affect investment and national saving. The natural rate of interest is often considered as a useful guide for monetary policy. By setting the real policy rate in line with the natural rate, monetary policy tends to accommodate fluctuations in the economy so as to bring output in line with potential output, at least in simple models. Furthermore, to the extent that inflation depends on the output gap, i.e., the gap between output and its potential, monetary policy could in principle stabilize output around its potential as well as inflation.

In their paper “Monetary Policy, the Financial Cycle, and Ultra-Low Interest Rates,” Mikael Juselius, Claudio Borio, Piti Disyatat, and Mathias Drehmann (henceforth JBDD) criticize this approach

*The views expressed herein are solely those of the author and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System.

as being too narrowly focused on inflation and the output gap, and as omitting, in their view, a key driver of economic fluctuations: the financial cycle. The paper's contributions are threefold. First, it proposes a thought-provoking account of the U.S. economy's evolution over the past three decades, in which monetary policy has a first-order impact on the financial cycle. Fluctuations in financial variables such as leverage and the debt service burden cause in turn inefficient output booms and busts. This results in large fluctuations in the output gap. Second, the authors revisit the measurement of the natural rate of interest. While r^* is, in popular analyses, the real rate that equates the output with the potential output, JBDD propose an alternative natural rate that accounts for the financial cycle. A key finding of the paper is that the estimate of the proposed natural rate tends to be higher than the more conventional r^* over the past three decades, and furthermore declines considerably less than conventional estimates during the recent Great Recession. The authors conclude from this that actual policy rates have been persistently lower than the natural rate that appropriately adjusts for financial cycles. Furthermore, they conclude that past loose monetary policy caused financial imbalances, output booms, and busts, which in turn led to the recent ultra-low rates. JBDD in turn recommend that monetary policy "lean against the wind" and respond systematically to short-term fluctuations in the financial cycle. They specifically propose a Taylor-type interest rate rule that responds to financial indicators, and argue, using a simple counterfactual experiment, that following such a policy rule would have improved economic and financial stability relative to the actual policy.

The paper raises numerous important questions: in particular, to what extent should policymakers worry about the financial cycle, above and beyond the mandated macroeconomic objectives? This paper provides a valuable contribution to this debate. It is very thought provoking and states important policy implications. The paper deserves attention and careful reading. The analysis is well executed, and the underlying working paper (JBDD 2016) contains important additions, numerous robustness checks, and appendixes. However, while the paper states strong conclusions and unequivocal policy implications, I am not convinced that the evidence provided in the paper supports these conclusions.

In the rest of my comments, I focus on the authors' characterization of the financial cycle, on the interpretation of the "finance-neutral" r^*, and argue that there are good reasons for the policy rate *not* to track the authors' proposed "finance-neutral" r^*. In the end, while I applaud the authors' valuable efforts to merge financial cycle considerations with a macroeconomic and monetary model, I remain unconvinced by the paper's main claim, namely, that policy should have tracked more closely the authors' "finance-neutral" r^*, which is on average higher and has declined less than other popular estimates.

2. The Financial Cycle

To determine the finance-neutral natural rate of interest, the authors augment the popular Laubach-Williams (2003) model with two variables that they argue characterize well the financial cycle. These variables, initially proposed in Juselius and Drehmann (2015), are the *leverage gap* and the *debt service gap*, each expressed in deviations from a respective long-run cointegrating relationship. The first relationship assumes that the credit-to-GDP ratio is in the long run directly proportional to the real price of assets.[1] The second relationship assumes that the credit-to-GDP ratio moves in the long run one-for-one with the average lending rate on the outstanding stock of credit, which depends on current and past monetary policy. The assumption of these two cointegrating relationships implies that any gap or deviation from these long-run relationships must reverse over time. For instance, JBDD argue that if leverage is below its long-run value (a negative leverage gap), which could be caused by unusually high asset prices, then credit would subsequently tend to rise faster than GDP to restore the leverage back to its "normal" level. A "financial equilibrium" obtains when both leverage and debt service gaps are closed.

While the assumption of cointegrating relationships may be reasonable in many circumstances, it is not innocuous. It does exclude

[1] Juselius and Drehmann (2015) justify this by assuming (i) that leverage (LEV_t) defined as the ratio of nominal credit (CR_t) to the nominal value of assets (A_t) is constant in the long run, and (ii) that the nominal value of assets divided by the nominal price of assets is proportional to real GDP.

permanent changes to a "normal" leverage level (say as a result of securitization, for instance) or to a "normal" debt service. The authors justify the long-run relationships not with theory, but by failing to reject the cointegration between those variables over the past three decades, and by arguing that the parameters involved have remained stable. I find this empirical evidence unconvincing. As Elliott (1998) demonstrated, standard hypothesis tests for cointegration depend heavily on the assumption of exact unit root in the model; for near unit root, these tests have low power to reject no cointegration. Moreover, the assumption of a unit root in interest rates is a priori questionable.

3. The Financial Cycle and the Macroeconomy

Assuming that we have appropriately characterized the financial cycle, the next question is how does it affect the macroeconomy? JBDD consider a variant of the model by Laubach and Williams (2003) composed of a dynamic IS equation of the form

$$\begin{aligned} y_t - y_t^* = \beta_3 \left(y_{t-1} - y_{t-1}^*\right) - \varphi_{31} \left(i_{t-1} - \pi_{t-1} - r_{t-1}^*\right) \\ - \varphi_{32}\widetilde{lev}_{t-1} + \vartheta_{3t}, \end{aligned}$$

where y_t is log output, y_t^* denotes log potential output, i_t is the nominal interest rate, π_t is inflation, r_t^* is the real natural rate of interest, $\widetilde{lev}_t$ denotes the leverage gap, and ϑ_{3t} is an exogenous shock; a Phillips curve

$$\pi_t - \pi^* = \beta_5 \left(\pi_{t-1} - \pi^*\right) + \varphi_5 \left(y_{t-1} - y_{t-1}^*\right) + \vartheta_{5t};$$

an equation determining the evolution of the natural rate of interest

$$r_t^* = \beta_6 r_{t-1}^* + (1 - \beta_6) \left(z_t + \rho^{-1} 4\Delta y_t^*\right) + \vartheta_{6t};$$

and exogenous processes describing the evolution of potential output y_t^* and z_t. This model is reminiscent of the simple New Keynesian model (e.g., Woodford 2003; Galí 2008). However, while the New Keynesian model results from agents' explicit optimization problems, which yield forward-looking relationships, the present model assumes backward-looking relationships that are harder to justify on theoretical grounds. That said, the model expands on the popular

Laubach-Williams (2003) model by adding the leverage gap to the IS equation.

The authors explain that the leverage gap is meant to loosely capture the strength of the credit constraint. A positive leverage gap implies that the credit-to-GDP ratio is high relative to asset prices. Given that the leverage gap is assumed to return to its initial steady state, going forward credit growth is expected to be lower than GDP growth. The authors argue that, somehow, this is meant to depress current GDP ($\varphi_{32} > 0$). While the authors take comfort in estimating a strong negative correlation between their leverage gap variable $\widetilde{lev}_t$ and the output gap $y_t - y_t^*$, I remain skeptical about the setup in the absence of a solid foundation and clear mechanisms justifying their equations.

3.1 On the Dangers of Backward-Looking Models

I have little doubt that the model provides a decent fit of the data. I am, however, concerned that the seemingly benign assumption of backward-looking relationships may have important implications for the interpretation of the results and hence the policy recommendations. This affects both the assessment of the effects of interest rate changes on economic activity and the influence of financial cycles on the economy.

First, estimates of the parameters entering the IS equation suggest that the leverage gap has a much greater impact on the output gap than interest rates: while the parameters estimates $\varphi_{31} = 0.037$ and $\varphi_{32} = 0.045$ are roughly the same, the leverage gap fluctuates considerably more than the interest rate in relation to the natural rate. Hence, according to the authors, the leverage gap "is one of the main channels through which monetary policy can influence the real economy." Furthermore, they find that the real interest rate is not a very effective tool for controlling the leverage gap itself. Having obtained that result, it is understandable that the authors conclude that a primary focus of monetary policy should be to stabilize the financial cycle instead of attempting to close the output gap directly. However, that argument relies crucially on the finding that interest rates have relatively little impact on output. That result is, however, in contrast with typical estimates of medium-scale dynamic stochastic general equilibrium models (e.g., Smets and Wouters 2007;

Del Negro, Giannoni, and Schorfheide 2015), or studies based on survey of the consumer expectations (Crump et al., 2015), which obtain considerably larger estimates of the intertemporal elasticity of substitution.

One key difference, as mentioned above, is that the proposed IS equation is backward looking. This makes it difficult to infer a structural estimate of the response of output to interest rate changes. To see why, imagine that interest rates have a large effect on output and the central bank sets its policy rate so as to stabilize the output gap next period, given all information available. If the central bank is successful, the interest rate movements should offset any shock that is about to move the output gap one way or the other. As a result, we should observe almost no correlation between the current policy rate and the next period's output gap. In such an environment, one would likely estimate a very low value for φ_{31}, even though interest rates have by assumption a large impact on output: the coefficient φ_{31} would confound both the impact of interest rate on output and the policy response. To address this problem, a structural model that disentangles the two forces is needed.

The second concern with the backward-looking specification refers to the link between the financial cycle and the macroeconomy. The financial cycle involves a potentially powerful feedback loop between the leverage gap, the debt service gap, asset prices, interest rates, and economic activity. It is characterized in the paper by the dynamic equations (8), (9), and (10). While the paper doesn't elaborate on this financial cycle, the underlying working paper (JBDD 2016) and Juselius and Drehmann (2015) illustrate it in more details. They argue that starting from a negative leverage gap (e.g., due to relatively high asset prices, as in the late 1980s or mid-2000s), credit would necessarily grow faster than output. As the credit-to-GDP ratio increases, so does the debt service gap. That ends up slowing down output and asset prices, eventually causing a bust and a potentially drawn-out recession. While these variables are certainly correlated, a natural question is which way does the causality go? The authors assume a priori that the financial cycle causes economic fluctuations. This leads naturally to the policy implications listed in their conclusion. The authors' preferred story has, however, the peculiar implication that leverage predicts future movements in asset prices, which are often thought to be purely forward looking.

An alternative view of the co-movement of financial and economic variables is that the financial cycle reflects economic fluctuations, whereby asset prices and credit are based on the expected future economic conditions. The causality would thus run in the other direction. It would be useful in future work to provide more evidence about the direction of causality.

4. Do "Prevailing" Models Neglect the Financial Cycle?

The authors argue that prevailing models neglect the financial cycle, and thus "present an alternative view of the natural rate, in which financial factors also play a role." It is true that in the simplest New Keynesian model—composed of a dynamic IS equation, a Phillips-curve relationship, and a monetary policy rule—there are no explicit "financial cycle" variables, and the natural rate of interest and of output are functions of real disturbances only. However, New Keynesian models with financial frictions do exist as well (e.g., Bernanke, Gertler, and Gilchrist 1999; Christiano, Motto, and Rostagno 2003, 2014; Cúrdia and Woodford 2009; Del Negro, Giannoni, and Schorfheide 2015; Del Negro et al. 2017, among many others). In such models, entrepreneurs' leverage is a key state variable which affects the spread between the nominal return on capital (the borrowing rate) and the policy rate (the deposit rate), along with shocks to the entrepreneurs' idiosyncratic productivity. In turn, aggregate output depends on the policy rate and credit spread, hence on leverage. But in those models, fluctuations in leverage need not be unsustainable; they can be efficient. It is thus not clear a priori that they need to be stabilized. For example, Del Negro et al. (2017) estimate a medium-scale DSGE model with financial frictions, and in which Treasury securities are valued not only for their pecuniary return but also for their liquidity and safety attributes, so that they carry a "convenience yield," along the lines of Krishnamurthy and Vissing-Jorgensen (2012) and Caballero, Farhi, and Gourinchas (2017). Interestingly, while the resulting natural rate of interest can fluctuate sharply in the short run, the implied medium-term natural rate (i.e., the five-year forward natural rate) resembles closely the estimate obtained by Laubach and Williams, over the past three decades.

Although these models incorporate some financial features, they remain very stylized. In particular, they typically do not capture the non-linear effects of high leverage on the economy's vulnerability to adverse shocks, which may trigger financial and economic crises. Model improvements are certainly still badly needed in this area. Recent work by Adrian and Duarte (2017) offers an interesting path forward.

5. On the Interpretation of r*

When one considers financial frictions, several interest rates come into play, and a spread separates the rate of interest charged to borrowers and the one paid to depositors. Furthermore, rising financial frictions tend to push the borrowing rate up while they tend to depress the deposit rate. Which rate does the natural rate of interest correspond to? To the extent that we are interested in an interest rate relevant for monetary policy, the relevant natural rate is the rate on a short-term riskless, liquid security.

The estimated model yields estimates of the latent variables r_t^* and y_t^* which the authors call "the 'finance-neutral' natural interest rate and potential output—in the sense that the estimates control for the influence of financial factors." JBDD then compare their estimate of r_t^* to that of Laubach and Williams (2003) and find that the "finance-neutral" rate is generally higher than the one estimated by Laubach and Williams (2003) and has also declined less during the Great Recession. The authors imply that the real policy rate should have been higher than implied by Laubach and Williams.

Taking JBDD's model literally, is the authors' "finance-neutral" rate really the rate that policy should track? I don't think so. Looking at the IS equation and the Phillips curve again, and abstracting from exogenous shocks ϑ_{3t} and ϑ_{5t} for simplicity, it is apparent that the central bank could in principle perfectly stabilize both the output gap and inflation around its target, by setting the policy rate to

$$i_t = \pi_t + \left(r_t^* - \frac{\varphi_{32}}{\varphi_{31}} \widetilde{lev}_t \right)$$

at all dates. As a result, in this model, policy should not track the "finance-neutral" r_t^*, but rather the natural rate adjusted by the leverage gap. With a strongly positive leverage gap in the midst of the Great Recession (see figure 2 in JBDD's paper), the model suggests that it was appropriate for policy to lower the interest rate well below the authors' estimate of r^* to stabilize the output gap and inflation. In practice, of course, the effective lower bound on nominal interest rates is likely to prevent the policy rate from tracking $r_t^* - \frac{\varphi_{32}}{\varphi_{31}}\widetilde{lev}_t$, especially given the large fluctuations in the leverage gap. It would be interesting then to compare the adjusted r^* to that obtained by Laubach and Williams (2003). The key point, though, is that JBDD's estimate of r^* is not a reference rate that helps close the output gap and the inflation gap.

What does JBDD's r^* refer to? Equation (8) in the paper reveals that r^* corresponds to the real rate at which policy would have no direct impact on the leverage gap. By tracking fluctuations in the estimated r^*, policy would, however, still indirectly contribute to fluctuations in the leverage gap, as the implied changes in the policy rate would affect the *debt service ratio gap* (see equations (9) and (10) in the paper), which in turn have an effect on the leverage gap.

6. Conclusion

In conclusion, this is a very interesting and thought-provoking paper on a key monetary policy issue. It presents an interesting framework that allows the authors to perform a valuable analysis of the complex interactions between the financial cycle and macroeconomic fluctuations. The paper concludes with strong policy implications. While the motivation for the model is based on observed empirical co-movements between financial and macroeconomic variables, I am concerned that some of the apparently innocuous assumptions affect importantly the results. In particular, I suspect that the conclusions rely heavily on the backward-looking nature of the model. In future work, I would encourage the authors to develop a theoretical justification for their setup or to explore the robustness of their conclusions to other models. I also remain unconvinced that the natural rate estimated by the authors provides a good guide for policy. That said, this paper provides an important contribution to the literature on macroeconomic and financial stabilization, and I expect it to remain a valuable reference.

References

Adrian, T., and F. Duarte. 2017. "Financial Vulnerability and Monetary Policy." Staff Paper No. 804, Federal Reserve Bank of New York.

Bernanke, B. S., M. Gertler, and S. Gilchrist. 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." In *Handbook of Macroeconomics*, Vol. 1C, ed. J. B. Taylor and M. Woodford, 1341–93 (chapter 21). Amsterdam: North-Holland.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas. 2017. "The Safe Assets Shortage Conundrum." *Journal of Economic Perspectives* 31 (3): 29–46.

Christiano, L. J., R. Motto, and M. Rostagno. 2003. "The Great Depression and the Friedman–Schwartz Hypothesis." *Journal of Money, Credit and Banking* 35 (6, part 2): 1119–97.

———. 2014. "Risk Shocks." *American Economic Review* 104 (1): 27–65.

Crump, R. K., S. Eusepi, A. Tambalotti, and G. Topa. 2015. "Subjective Intertemporal Substitution." Staff Paper No. 734, Federal Reserve Bank of New York.

Cúrdia, V., and M. Woodford. 2016. "Credit Frictions and Optimal Monetary Policy." *Journal of Monetary Economics* 84 (December): 30–65.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti. 2017. "Safety, Liquidity, and the Natural Rate of Interest." *Brookings Papers on Economic Activity* (Spring).

Del Negro, M., M. P. Giannoni, and F. Schorfheide. 2015. "Inflation in the Great Recession and New Keynesian Models." *American Economic Journal: Macroeconomics* 7 (1): 168–96.

Elliott, G. 1998. "On the Robustness of Cointegration Methods when Regressors Almost Have Unit Roots." *Econometrica* 66 (1): 149–58.

Galí, J. 2008. *Monetary Policy, Inflation, and the Business Cycle.* Princeton, NJ: Princeton University Press.

Holston, K., T. Laubach, and J. C. Williams. 2017. "Measuring the Natural Rate of Interest: International Trends and Determinants." In *NBER International Seminar on Macroeconomics 2016*, ed. M. B. Devereux. National Bureau of Economic Research, Inc.

Juselius, M., C. Borio, P. Disyatat, and M. Drehmann. 2016. "Monetary Policy, the Financial Cycle and Ultra-Low Interest Rates." BIS Working Paper No. 569.

Juselius, M., and M. Drehmann. 2015. "Leverage Dynamics and the Real Burden of Debt." BIS Working Paper No. 501.

Krishnamurthy, A., and A. Vissing-Jorgensen. 2012. "The Aggregate Demand for Treasury Debt." *Journal of Political Economy* 120 (2): 233–67.

Laubach, T., and J. C. Williams. 2003. "Measuring the Natural Rate of Interest." *Review of Economics and Statistics* 85 (4): 1063–70.

Smets, F., and R. Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586–606.

Woodford, M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton, NJ: Princeton University Press.

Low Interest Rates: Causes and Consequences*

Robert E. Hall
Hoover Institution and Department of Economics,
Stanford University
National Bureau of Economic Research

World interest rates have been declining for several decades. In a general equilibrium setting, the interest rate is determined by the interaction of a number of types of behavior: the policy of the central bank, investment in productive assets, the choice between current and future consumption, and the responses of wealth holders to risk. Central banks devote consider effort to determining equilibrium real rates, around which they set their policy rates, though measuring the equilibrium rate is challenging. The real interest rate is also connected to the marginal product of capital, though the connection is loose. Similarly, the real interest rate is connected to consumption growth through a Euler equation, but again many other influences enter the relationship between the two variables. Finally, the idea of the "global saving glut" suggests that the rise of income in countries with high propensities to save may be a factor in the decline in real rates. That idea receives support in a simple model of global financial equilibrium between countries with risk tolerance (the United States) and ones with high risk aversion (China).

JEL Codes: E21, E22, E43, E52.

Low world interest rates have stimulated new interest in the determination of the safe real rate. As a threshold matter, Rachel and Smith's figure 1 (this issue) and Juselius et al.'s figure 1 (this issue) document the pronounced downward trend of world real interest rates since the 1980s. For the purposes of this commentary, I take

*This research was supported by the Hoover Institution. Complete backup for all of the calculations is available from my website, http://www.stanford.edu/∼rehall. Author contact: rehall@stanford.edu; stanford.edu/∼rehall.

the real rate to be the yield net of inflation of safe government debt of maturity around one to two years. Thus I abstract from liquidity effects at the short end of the yield curve and from issues related to the slope of the yield curve.

Structural relations governing the real interest rate include its relation to

- the central bank's payment on reserves and the extent of saturation of the financial system in reserves
- the marginal product of capital
- the rate of consumption growth (through the Euler equation)
- the terms of trade between risk-tolerant and risk-averse investors

In a complete macro model one or more equations would describe each of these structural relations. It would not be possible to divide up responsibility among them for the overall decline in the real rate. One can fashion a set of highly simplified models, each containing only one or two of the structural relations. For example, Krugman (1998) considers an economy with no capital and no uncertainty to focus on monetary policy and consumption growth and illuminate issues of the zero lower bound. But a set of models along those lines would not result in an additive breakdown of the sources of the decline in the real interest rate.

1. Monetary Policy and the Real Interest Rate

Traditional monetary policy kept the interest paid on reserves at zero nominal and manipulated the quantity of reserves. Explaining how the central bank influenced interest rates involved consideration of the liquidity value of scarce reserves. Today, all major central banks have saturated their financial systems with reserves, so the liquidity value is zero, and the central banks execute monetary policy exclusively by manipulation of the payment made to reserve holders (in the United States, a new kind of reserves, reverse repurchase agreements, play this role).

Powerful forces of arbitrage link the central bank's policy rate paid on reserves to similar short-term government obligations. The central bank thus controls short rates directly. But the fact of central bank control does not mean that we need look no further to understand the movements of short rates. For one thing, it is the behavior of real rates that matters and all central banks set nominal rates, though there would be no obstacle to direct setting of real rates. Hall and Reis (2016) discusses these topics in detail. Thus the behavior of inflation needs to be brought into the picture. More important, however, is that changing the policy rate has effects on output and employment relatively quickly and on inflation, with a longer lag, according to most views.

As a result of the influence of the central bank's policy rate on the other key macro variables, the other structural relations listed above come into play in the central bank's choice of the policy rate. Only the most naive observer thinks that the central bank can pick its policy rate by unilateral whim. Friedman (1968), following Wicksell, set forth a framework that remains influential fifty years later: There is a level of the real interest rate, r^*, the *natural rate*, with the property that it is infeasible for the central bank to run a monetary policy that results in a real rate permanently above or below the natural rate. Thus many discussions of the behavior of the real rate focus on quantifying r^*, generally as a quantity that varies over time. Since 1980, it has had a downward trend.

The foundations of the hypothesis that r_t^* is a cognizable feature of the economy are weak, in my opinion—see Hall (2005). It takes an economic or statistical model to extract r_t^* from data on r_t and other variables. The results are model specific. Laubach and Williams (2003) is the canon of this literature. Notwithstanding my doubts about the foundations, these authors' results seem completely reasonable. Juselius et al. (this issue) refine the canon. The middle of their figure 6 shows the real rate, which is volatile and cyclical. The Laubach-Williams natural rate is a plausibly smoothed version of the actual real rate. As Friedman's analysis predicted, the actual real rate exceeds its natural level in booms and falls below in busts. The natural rate of Juselius et al. has higher volatility and, surprisingly, a higher level. Friedman's analysis suggested fairly persuasively that the real rate should deviate above about as much as below the natural rate, but the new construction has almost all of the deviations below.

Figure 1. Spread between the Return to Capital and the Safe Real Interest Rate

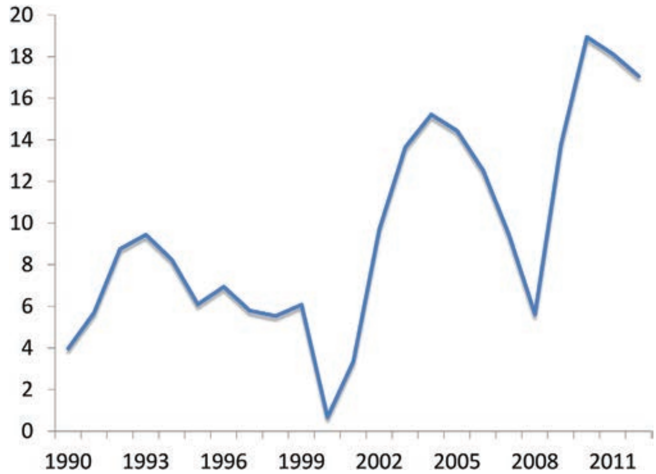

2. The Marginal Product of Capital and the Return to Capital

In an economy without uncertainty, the return to capital is linked to the marginal product by the rental price of capital. Provided the rental price includes the fluctuations in Tobin's q—the ratio of the value of installed capital to the acquisition price of capital—arbitrage should equate the marginal product of capital to the rental price. To put it differently, if the rate of return is calculated from data that accounts for q, the rate of return will track the interest rate (measured over the same interval) period by period. With uncertainty, the rate of return will include a risk premium, which may vary over time. The recent macro literature has studied financial frictions that interpose between wealth holders and businesses seeking to attract wealth to form business capital.

Figure 1 shows the spread between the calculated return to capital and the one-year safe real interest rate, from Hall (2015). Note that the spread is remarkably volatile, upward trending, and high except in recessions. Gomme, Ravikumar, and Rupert (2015) have made similar calculations. The notion that there is a tight connection between the safe interest rate and the return to capital receives

Figure 2. U.S. Real Rate and Consumption Growth

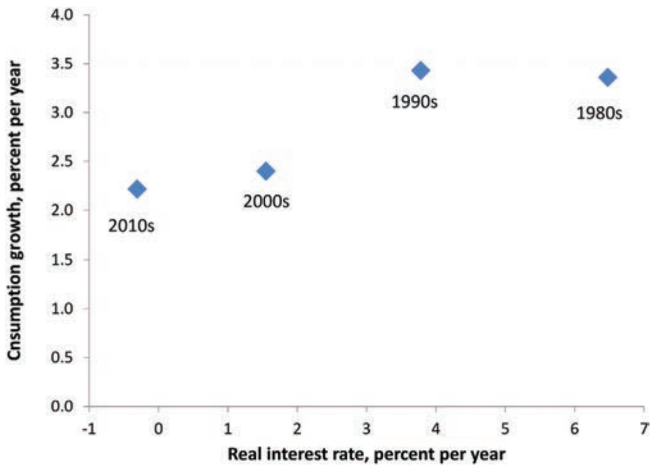

little support from this evidence. Rather, there is apparently large scope for variations over time in risk premiums, financial frictions, and other sources of the wedge between the earnings of capital and the risk-free cost of borrowing. These variations are almost certainly endogenous.

3. Consumption Growth and the Interest Rate

Many macro models, including the New Keynesian models that have proliferated at central banks, contain an upward-sloping structural relation between expected consumption growth and the real interest rate—Rachel and Smith's equation (1) describes the Euler equation reflecting this relation. The logic is that a higher real interest rate makes future consumption cheaper than current consumption, so households consume less currently and more in the future. To put it another way, higher growth rates should have correspondingly higher real interest rates. Figure 2 shows that this proposition is somewhat true in U.S. data averaged over decades.

The proposition encounters some serious obstacles. First, Carroll and Summers (1991) observed that across countries that can trade goods and financial claims, all countries should have the same rate of growth of consumption, in accord with the worldwide real interest rate, irrespective of their rates of growth of income. Countries with high expected income growth should borrow from slower-growing countries and gradually pay the debt off as growth occurs. In fact, the evidence shows that consumption growth is tightly linked to income growth across countries. And growth rates differ markedly across countries, with the highest growth in recent decades in east and south Asia.

Second, a household does not have a single Euler equation, but rather a different one for each asset. Hansen and Singleton (1983) is the classic citation on this point. There is nothing special about the safe real interest rate. Their paper showed that the data rejected the hypothesis that households satisfied all of the Euler equations.

Third, data on household financial holdings make it clear that households with collectively an important fraction of total income face binding constraints on borrowing. They would like to obey the Euler-equation model but cannot commit to repaying the debt that they would incur if they did. They obey a related model where a shadow borrowing rate, higher than the measured one, tracks consumption growth.

I conclude that research on consumption choices has a far richer view than the one expressed in the simple interest-only Euler equation.

4. The Role of the Interest Rate in an Economy where Risk-Tolerant Investors Insure Risk-Averse Ones by Borrowing from Them

Hall (2016) demonstrates the theoretical and practical importance of trade among heterogeneous investors. In effect, the risk-tolerant investors insure the risk-averse ones. Debt has a key role in this risk-motivated trade. By borrowing from the risk averse, the risk-tolerant investors provide the risk averse with protection against future random shocks, because the payoff of the debt is unaffected by the shocks (provided no default occurs). The interest rate on the debt describes the terms of the risk trade. If the risk tolerant have high resources relative to the risk averse, collectively, the risk averse command a good deal—they receive a high rate of interest on the funds they loan to the risk tolerant. But if there is an upward trend in the resources of the risk averse, the deal shifts disadvantageously away from the risk averse—they earn less and less interest on the funds they lend. The paper shows that China behaves risk aversely, lending large volumes of funds to western Europe and the United States. But the Chinese resource base—measured by GDP—is growing faster than the resource base of the risk-tolerant borrowers. Hence the world real interest rate is declining on account of the differential growth.

The model backing up this analysis is rigged to avoid the other issues discussed earlier in this commentary. There is no central bank intervening in the world financial market. There is no capital, so no issue of the relation of the marginal product of capital to the interest rate. Resources are growing at the rate of zero among the risk averse and the risk tolerant, so there are no issues of growth affecting the interest rate. The model embodies standard ideas from financial markets, including the hypothesis that investors attribute a small but positive probability that a truly bad event will occur and the hypothesis that the risk-averse investors place a somewhat higher probability on that event.

My paper pursues the ideas in Bernanke et al. (2011) that there is a "global savings glut" and in Gourinchas, Rey, and Govillot (2010) and Caballero and Farhi (2016) that low real interest rates are the result of a "shortage" of safe assets. The paper derives results along those lines from the equilibrium of an Arrow-Debreu economy with complete capital markets. In place of gluts and shortages, the model hypothesizes changes over time in the resources held by the risk tolerant in relation to those held by the risk averse.

Figure 3 shows how the safe real interest rate in the model declines as the fraction of resources held by the risk tolerant declines. The decline is similar to the decline that actually occurred from 1990 to the present, with real rates at or below zero. The risk-tolerant investors in the model have modestly lower coefficients of relative risk aversion and believe that the probability of bad conditions is modestly lower, compared with the risk-averse investors.

The conclusion of the model is that heterogeneity coupled with a shift in relative resources toward the risk-averse investors can explain observed changes in the real interest rate without bringing in the declining growth rate or rising financial frictions. The paper makes no claim that the other forces are not actually influential, however. Fundamental to the success of the model is its hypothesis that both

Figure 3. As the Fraction of Resources in the Hands of the Risk Tolerant Declines, the Interest Rate Falls

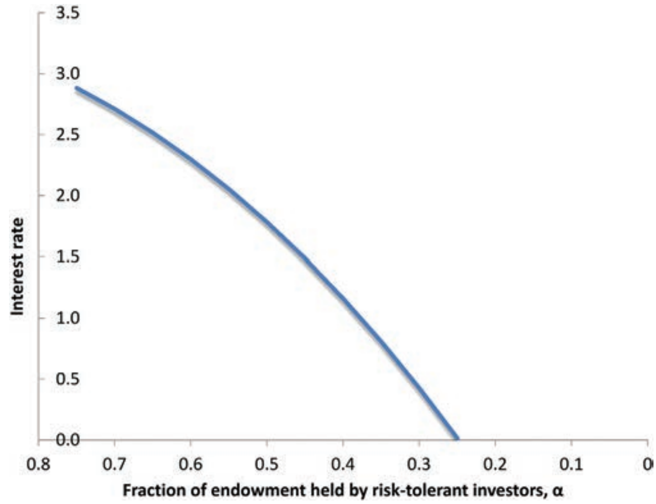

types of investors behave as if they assigned small but important probabilities to a substantial negative shock, worse than has actually occurred since the Great Depression. In this respect, the model follows the trend in recent financial economics, which finds, for example, that such beliefs about rare disasters are the most plausible way to explain the equity premium.

One of the manifestations of heterogeneity in investors' risk aversion is across countries. Investors in some countries, notably the United States, collectively take on risk from other parts of the world by maintaining positive net positions in foreign equity and negative net positions in debt—in effect, these countries borrow from the risk-averse countries and use the proceeds to buy foreign equity. Thus the United States is like a leveraged hedge fund. Countries can be divided into three groups: (i) those that absorb risk by borrowing in the global debt market and buying foreign equity, (ii) those that shed risk by lending to the risk absorbers and letting those countries take on the risk of their own equity, and (iii) those whose risk preferences are in the middle and choose not to absorb or shed risk and those whose financial markets are undeveloped and do not participate in global financial markets.

Figure 4. Countries that Absorb Risk by Holding Positive Amounts of Net Foreign Equity or by Borrowing from Foreign Lenders

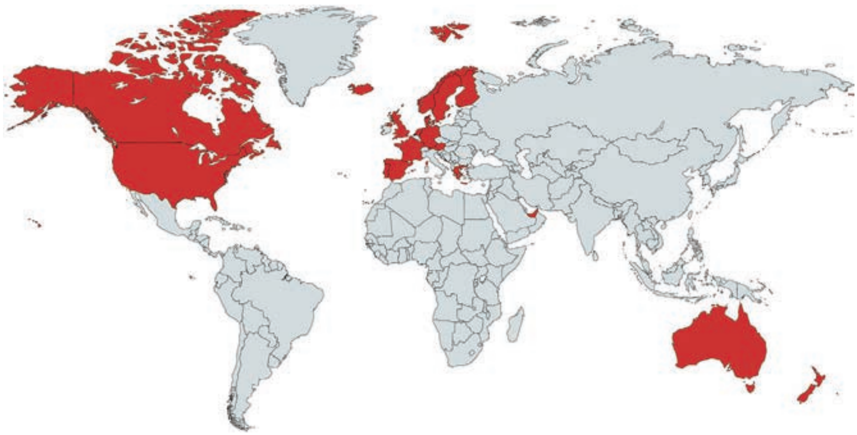

Note: Risk-absorbing countries are shown by dark shading. Created with mapchart.net.

Figure 4 shows the countries that absorb risk. They are the advanced countries of western Europe and the countries scattered around the globe that fell under the influence of those countries and became advanced themselves. There appears to be a negative correlation between risk aversion and income per person, as the risk absorbers are all high-income countries. By far the largest absorber of risk is the United States.

Figure 5 shows the countries that shed risk. Most are lower income. China is by far the largest of the shedders. China holds large amounts of dollar debt claims on the United States, with recent growth in its euro debt claims on western Europe. One high-income country, Japan, is a major risk shedder. The United States and other risk absorbers hold positive net amounts of foreign equity.

Figure 6 shows the growth of risk absorption by the United States. The upper line shows U.S. net borrowing in the debt market and the lower line net U.S. holdings of foreign equity. The upward path in debt began in the mid-1980s and the upward path of equity in the 1990s. Debt continued to rise through 2011 (the last year for which I have data) while equity fell slightly after the 2008 financial

Figure 5. Countries that Shed Risk by Holding Negative Amounts of Net Foreign Equity or by Lending Positive Amounts to Foreign Borrowers

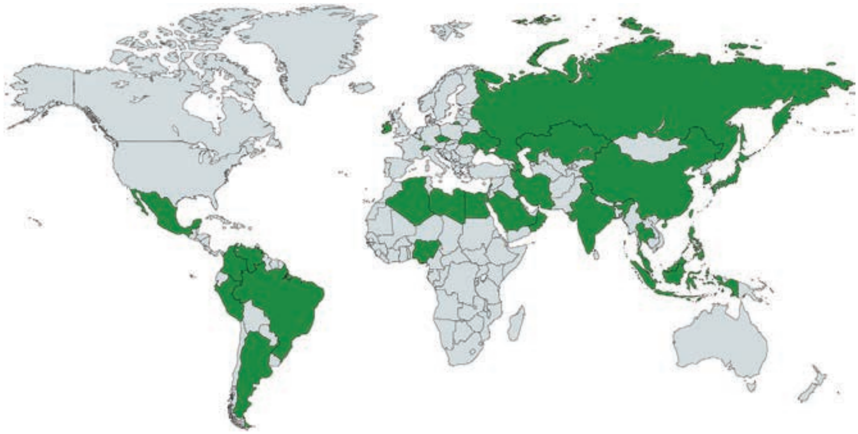

Note: Risk-shedding countries are shown by dark shading. Created with mapchart.net.

Figure 6. Risk Absorption by the United States, 1970–2011

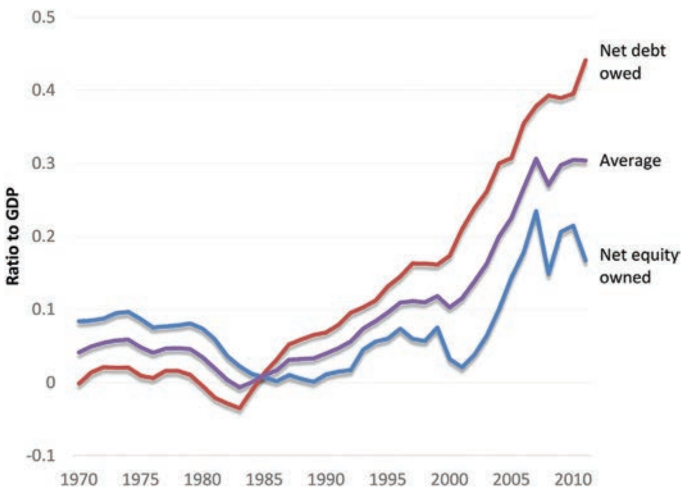

Figure 7. Risk Shedding by China, 1981–2011

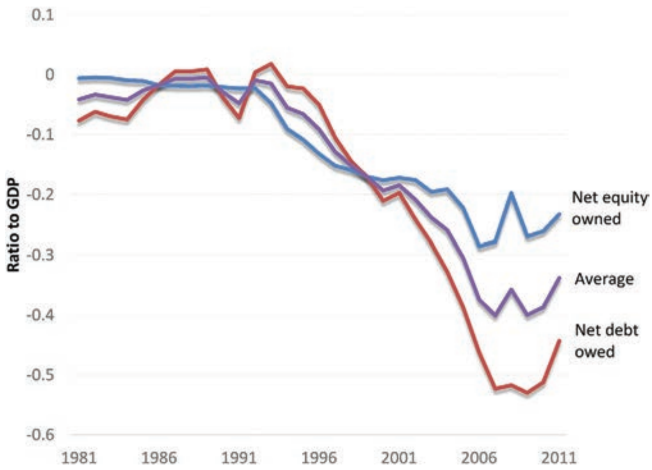

crisis. The average of the two measures—taken as an overall measure of risk absorption—rose from the 1980s and reached a plateau of 0.3 years of GDP.

Figure 7 shows similar data for China starting in 1981—in earlier years, China was effectively walled off from the global economy. Starting in the early 1990s, China shed risk aggressively, reaching the point just before the crisis of the average of foreign debt owned and net foreign holdings of Chinese equity claims equal to 0.4 years of GDP. Following the crisis, Chinese risk shedding has remained at that level but has not grown.

Risk splitting occurs within the United States in large volumes as well. Table 1 shows decade averages of a variety of financial institutions that hold risky financial positions funded in part by debt—held by risk-averse investors such as pension funds—and by correspondingly riskier equity held by risk-tolerant investors such as high-wealth households. Government debt is a prominent part of the risk splitting. In the case of government, the taxpayers make up the risk-tolerant side—the marginal taxpayer with substantially higher than average wealth takes on magnified risk by insuring the holders of government debt. On the private side, numerous types of financial institutions and securities have the effect of splitting risk between a tranche of low-risk debt and high-risk residual equity claims.

Table 1. Examples of the Scale of Risk-Splitting Institutions

	Government			Private				
Decade	**Consolidated Government Debt**	**GSE Debt**	**GSE Guaranteed Debt**	**Private Equity Funds**	**Securiti-zations**	**Non-financial Corporate Debt**	**Repos**	**Non-mortgage Household Debt**
1980s	0.469	0.061	0.091	—	0.012	0.163	0.103	0.186
1990s	0.611	0.101	0.204	—	0.086	0.211	0.166	0.204
2000s	0.574	0.203	0.293	0.058	0.233	0.238	0.237	0.239
2010s	0.936	0.126	0.347	0.140	0.109	0.275	0.221	0.251

Figure 8. Scale of Risk-Splitting Institutions Relative to GDP

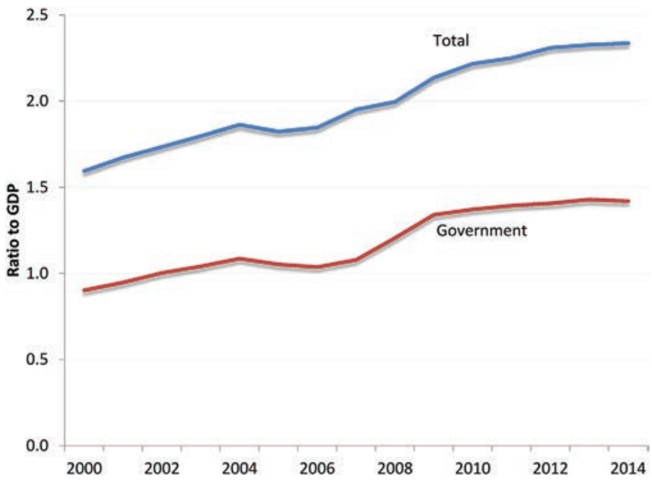

Private equity is a rapidly growing example of this type of financial arrangement. Securitizations with overcollateralized debtlike securities held by or on behalf of risk-averse investors and residual equity claims held by risk-tolerant investors grew rapidly until the crisis but have shrunk since then. Repurchase agreements split risk by overcollateralization to the extent of the repo haircut. These too have shrunk relative to GDP since the crisis.

Figure 8 shows the generally upward trend of the volume of risk splitting in the United States, stated relative to GDP. Both government and non-government contributions have risen, with some moderation after the crisis.

5. Concluding Remarks

Prior to the financial crisis in 2008, risk splitting grew steadily, as revealed in data on both international and domestic financial positions. Safe real interest rates declined in parallel. The crisis resulted in a downward jump in real rates corresponding to the fall in nominal short rates to essentially zero soon after the crisis struck. The corresponding real rate was between –1 percent and –2 percent. Real rates have risen in the United States recently, as nominal rates have become positive and inflation has risen close to the Federal Reserve's target of 2 percent, but real rates in other markets remain as negative as ever in the eight years since the crisis. Because the crisis hit GDP and asset value harder in advanced countries than in others, especially China, the influence studied in my analysis may explain some part of the drop in the global safe real short rate. In addition, the crisis may have raised investors' beliefs about the probability of adverse events in the future, as in Kozlowski, Veldkamp, and Venkateswaran (2015). According to the principles considered here, the safe real rate would fall if the disaster probability rose more for the risk-averse investors than for the risk tolerant.

I emphasize again that heterogeneity in risk aversion is only one of the factors entering a full explanation of the behavior of real rates over recent decades. Expansionary monetary policy, rising financial frictions, and slowing consumption growth need to be brought into a full analysis.

References

Bernanke, B. S., C. Bertaut, L. Pounder DeMarco, and S. Kamin. 2011. "International Capital Flows and the Returns to Safe Assets in the United States, 2003–2007." International Finance Discussion Paper No. 1014, Board of Governors of the Federal Reserve System (February).

Caballero, R. J., and E. Farhi. 2016. "The Safety Trap." March. Harvard University, Department of Economics.

Carroll, C. D., and L. H. Summers. 1991. "Consumption Growth Parallels Income Growth: Some New Evidence." In *National Saving and Economic Performance,* ed. B. D. Bernheim and J. B. Shovin, 305–48 (chapter 10). University of Chicago Press.

Friedman, M. 1968. "The Role of Monetary Policy." Presidential address delivered at the 80th Annual Meeting of the American Economics Association, Washington, DC, December 29, 1967. *American Economic Review* 58 (1): 1–15.

Gomme, P., B. Ravikumar, and P. Rupert. 2015. "Secular Stagnation and Returns on Capital." *Economic Synopses* (Federal Reserve Bank of St. Louis) (19): 1–3.

Gourinchas, P.-O., H. Rey, and N. Govillot. 2010. "Exorbitant Privilege and Exorbitant Duty." Discussion Paper No. 10-E-20, Institute for Monetary and Economic Studies, Bank of Japan.

Hall, R. E. 2005. "Separating the Business Cycle from Other Economic Fluctuations." In *The Greenspan Era: Lessons for the Future,* 133–79. Proceedings of a symposium sponsored by the Federal Reserve Bank of Kansas City, August 25–27.

———. 2015. "Quantifying the Lasting Harm to the U.S. Economy from the Financial Crisis." *NBER Macroeconomics Annual 2014*, Vol. 29, ed. J. A. Parker and M. Woodford, 71–128. University of Chicago Press.

———. 2016. "The Role of the Growth of Risk-Averse Wealth in the Decline of the Safe Real Interest Rate." Hoover Institution (November).

Hall, R. E., and R. Reis. 2016. "Achieving Price Stability by Manipulating the Central Bank's Payment on Reserves." NBER Working Paper No. 22761 (October).

Hansen, L. P., and K. J. Singleton. 1983. "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *Journal of Political Economy* 91 (2): 249–65.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran. 2015. "The Tail that Wags the Economy: Beliefs and Persistent Stagnation." NBER Working Paper No. 21719 (November).

Krugman, P. R. 1998. "It's Baaack: Japan's Slump and the Return of the Liquidity Trap." *Brookings Papers on Economic Activity* (2): 137–205.

Laubach, T., and J. C Williams. 2003. "Measuring the Natural Rate of Interest." *Review of Economics and Statistics* 85 (4): 1063–70.

Macroprudential Policy under Uncertainty*

Saleem Bahaj and Angus Foulis
Bank of England and Centre for Macroeconomics

We argue that uncertainty over the impact of macroprudential policy need not make a policymaker more cautious. Our starting point is the classic finding of Brainard that uncertainty over the impact of a policy instrument will make a policymaker less active. This result is challenged in a series of richer models designed to take into account the more complex reality faced by a macroprudential policymaker. We find that asymmetries in policy objectives, the presence of unquantifiable sources of risk, the ability to learn from policy, and private-sector uncertainty over policy objectives can all lead to more active policy.

JEL Codes: D81, E58.

1. Introduction

The macroprudential toolkit available to policymakers across several central banks is new and largely untested. For example, in the United Kingdom, the Bank of England's Financial Policy Committee (FPC) has, since the financial crisis, received powers to alter bank capital requirements and to place restrictions on the terms of household mortgages for macroprudential purposes. Neither of these policy tools has been used previously, so their impact and the Committee's reaction function remain unclear. Moreover, in contrast to

*We are grateful to the editor, Rafael Repullo, our discussant, Francois Gourio, and to participants at the 2016 Annual Conference of the *International Journal of Central Banking* for helpful comments and feedback. We would also like to further thank David Aikman, John Barrdear, Oliver Bush, Julia Giese, Rodrigo Guimaraes, Alexandre Kohlhas, Marc Hintersweiger, Sujit Kapadia, Frederic Malherbe, Roland Meeks, Tamarah Shakir, Rhiannon Sowerbutts, Misa Tanaka, and Martin Taylor. This paper contains the views of the authors and not necessarily those of the Bank of England, the MPC, the FPC, or the PRA board. The working paper version of this paper is Bahaj and Foulis (2016) and contains additional results that we reference in the text. Author contact: saleem.bahaj@bankofengland.co.uk; angus.foulis@bankofengland.co.uk.

monetary policy, where price stability can be judged against the rate of inflation, the objective of macroprudential policymakers, the stability of the financial system, is inherently unobservable. Thus macroprudential policymakers face a high degree of uncertainty over the impact and effectiveness of their tools and a target variable they cannot perfectly observe. In the face of this uncertainty, a prevalent view is that a cautious approach is warranted: if a policymaker is unsure what a tool does, she should use it gingerly. Indeed, this is a classic result from the literature on optimal policy under uncertainty as shown in Brainard (1967).

This paper takes the Brainard model as a starting point and asks: is the uncertainty faced by macroprudential policymakers sufficient to justify a cautious stance to macroprudential policy? The Brainard framework is stylized and static, and there are multiple reasons why a policymaker may want to overlook its conclusions. In this paper, we present the results from some simple extensions to this framework to illustrate how uncertainty could alter the behavior of policymakers. The analysis here is drawn from the existing literature, but our goal is to frame the issue of uncertainty in the macroprudential context.

As a starting point, and to fix ideas, we recast the Brainard model as a macroprudential policy problem where the policymaker attempts to stabilize the resilience of a financial system. In particular, we assume that the policymaker is trying to stabilize the level of financial stability denoted x about some target x^* through the use of a tool k (for example, a time-varying capital requirement such as the countercyclical capital buffer) that controls x imperfectly. The relationship between k and x is linear,

$$x = bk + u, \tag{1}$$

and is subject to two sorts of uncertainty. First, b, the parameter governing how k affects x, is uncertain with prior mean $\bar{b} > 0$ and variance σ_b^2. The failure of the policymaker to observe b perfectly could, for instance, reflect uncertainty over the impact of capital requirements on financial stability. Second, there is unobserved variation in the level of financial stability, u, which is independent of the policymaker's action and has prior mean 0 and variance σ_u^2. To simplify the exposition, we assume the two sources of uncertainty are uncorrelated in what follows. This is a standard assumption that model and shock uncertainty are not related; however, similar results do emerge in a more general setting.[1]

Policymakers should find an unstable financial system undesirable; however, an overly stable system may dampen economic activity and impose a burden on the financial intermediaries or consumers. To cite a cliche: policy should aim to avoid the stability of the graveyard. The value of x^* can, therefore, be thought of as the optimal level of financial stability, trading off a stable versus an active financial system. Similarly, adjusting k may also impose costs; for example, by forcing banks to pay the underwriting fees associated with equity issuance. Thus we assume the policymaker has the following objective:

$$W = -\frac{1}{2}\mathbb{E}((x - x^*)^2 + \lambda k^2). \tag{2}$$

The parameter $\lambda > 0$ captures the policymaker's view over the relative cost of stabilizing k versus x about their optimal levels.[2] Under these assumptions, it is straightforward to take the policymaker's first-order condition ($\mathbb{E}[b^2 k + bu - bx^* + \lambda k] = 0$) and show that the best choice of k under uncertainty (denoted k^u) is

$$k^u = \frac{\bar{b}x^*}{\bar{b}^2 + \sigma_b^2 + \lambda} < \frac{\bar{b}x^*}{\bar{b}^2 + \lambda} = k^c, \tag{3}$$

where k^c denotes the level of k the policymaker should choose if she faced no uncertainty. Policy is less active under uncertainty. Further, as uncertainty over the impact of policy, σ_b^2, increases, k^u falls. This means that the policymaker should use her tool "less" as uncertainty increases, and for any given level of uncertainty the policymaker should choose k at a lower level than if she was certain.

[1]The more general case of $Cov(b, u) \neq 0$ is discussed in the working paper version of this paper.

[2]Note that we are assuming that the policymaker has a symmetric objective. This is potentially unrealistic both for financial stability and the costs of changing k: low financial stability may be more worrisome than high financial stability, and cutting capital requirements may not impose much of a burden on the financial system relative to raising them. In section 2 we consider an asymmetric financial stability objective for the policymaker with low financial stability disproportionately costly.

The intuition for this result is simply that additional uncertainty over the tool is perceived to introduce additional volatility into the economy when it is used, which is undesirable from the policymaker's perspective. A policy instrument whose impact is uncertain should be used more sparingly. Further, when there is greater uncertainty, the instrument should be used less. In this static model we associate the degree of policy activism with the level of k chosen. A more natural interpretation of "activism" might be the responsiveness of policy to shocks. The model can support such an interpretation by considering the response of k to the desired level of financial stability x^*—greater uncertainty over the impact of policy will make the policy less responsive to the realized value of x^*. In section 5 we explicitly consider a model in which policy is set after observing shocks to financial stability.

An important feature of Brainard's analysis is that the form of uncertainty matters. Brainard's results are sometimes misleadingly cited as a general rule that a policymaker should do less in the face of uncertainty. However, note that σ_u^2 does not appear in k^u. Therefore, the second conclusion from this form of model is that being unsure over the state of the economy (for example, the inherent stability of the financial system) should not alter policymakers' behavior.

This paper argues that there are several types of uncertainty and multiple channels through which uncertainty can affect policymaking. The result that policy should be more cautious in the presence of uncertainty does not hold in general, particularly for specific examples that are relevant to macroprudential policy. As we shall see, if anything, the results speak to a more active policy stance in the face of uncertainty. This complements the need for policymakers to guard themselves against inaction bias. Financial stability risks are hard to measure (or unobservable), and actions to address them may have short-term costs making regulatory forbearance tempting. The lags associated with macroprudential policy instruments, both in terms of implementation and transmission, mean that there is the potential for policymakers to move too late to build resilience in the financial system ahead of crises.

To make these points, we consider several extensions to the Brainard model. Our first extension considers an asymmetric objective function for financial stability in which a crisis is disproportionately costly for the policymaker. If the policy tool is sufficiently effective on average, policy will become more active the greater this asymmetry. In a second extension, we recognize that macroprudential policy is concerned with rare events, the probability of which is difficult to quantify. In such a situation, the policymaker may wish to behave in a robust fashion, preparing for the worst-case scenario. This can also lead to more active policy. Our third extension considers a dynamic model which allows for learning. Using the tool today reduces uncertainty about its impact tomorrow, but may initially increase volatility. If the motivation to learn is sufficiently strong (i.e., the policymaker's discount rate is sufficiently low), optimal policymaking can become more active with greater uncertainty. In our final extension, we consider the interaction between private-sector uncertainty and the uncertainty facing the policymaker. In addition to directly affecting financial institutions, macroprudential actions may have a broader impact through signaling information about risks to financial stability. We show that this signaling channel will be less powerful when there is greater private-sector uncertainty about policy objectives. Consequently, when the private sector is more unsure about why the policymaker is acting, the policymaker will need to be more active in order to offset the diminished signaling power of the tool.

1.1 Related Literature

Brainard (1967) lays out the canonical case for higher uncertainty leading to diminished policy activism. Beyond this classic work, this paper has several links with the academic literature on policy under uncertainty. However, most of this previous work has focused on the monetary policy context; our contribution is to recast some of the findings in the context of macroprudential policy and discuss their relevance.

In terms of policy under fundamental (or Knightian) uncertainty, the approach of using robust control and min-max dates back to at least Wald (1950); a more detailed, modern treatment is available in Hansen and Sargent (2001, 2008). Barlevy (2009) offers additional exposition over the framework. Robust control and min-max based optimal policy problems generally deliver more aggressive policy action in response. However, Onatski and Stock (2002) show that this finding does not extend to all model perturbations.

On the topic of learning through policy actions to reduce uncertainty, several papers—for example, Orphanides and Williams (2007) and Wieland (2000, 2006)—show how the desire to learn can lead to policy being more aggressive in the face of uncertainty. An alternative strand of the literature (see, for example, Besley 2001, Landes 1998, and Mukand and Rodrik 2002) argues that cross-sectional variation in policymaking across countries or political parties should also be encouraged to foster innovation and allow for learning about the impact of alternative policies. In contrast, other authors have challenged the gains from actively introducing policy variation in order to learn. In a quantitative assessment, Svensson and Williams (2007) show that the benefits for experimentation can often appear very modest in a generic linear quadratic forward-looking setup allowing for model uncertainty. Cogley and Sargent (2005) also cast doubt on the benefits of setting policy with learning in mind and raise a more general point that the gains from more active policy for the purposes of learning could be limited if there are sufficient natural experiments.

The interaction of the uncertainty of the public and policymaker, alongside asymmetric information and the informational value of government policy, are discussed in the early contributions of King (1982) and Weiss (1982). More recently, Lorenzoni (2010) and Angeletos, Iovino, and La'O (2011, 2015) analyze how to optimally set policy in static models with dispersed information.

2. Asymmetric Objective Function

The model in the introduction has a *symmetric* financial stability loss function $(x - x^*)^2$, in which financial stability being above the target level x^* is equally as costly to the policymaker as financial stability being below target. Such objective functions have been frequently used for modeling inflation targeting, as high inflation and deflation can both be costly. However, it is not clear that this functional form is appropriate for financial stability: the losses associated with a financial crisis may be significantly greater than those imposed by having excessive stability. In this section we consider an extension to the Brainard model with an asymmetric loss function that captures this property.

Figure 1. Graphical Illustration of the Linex Function

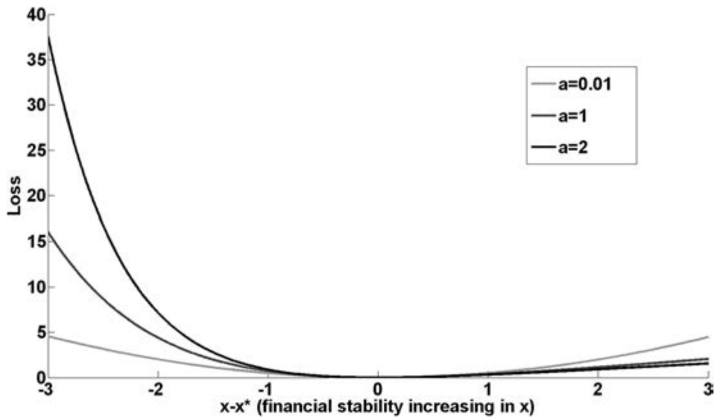

2.1 Modeling Framework

A tractable way to model an asymmetric loss function is to use the linex function (Varian 1975).[3] We consider the following objective function for the macroprudential policymaker:

$$W = -\frac{1}{a^2}\mathbb{E}\left[exp\left(a\left(x^* - x\right)\right) - a\left(x^* - x\right) - 1\right] - \frac{\lambda}{2}k^2, \qquad (4)$$

where $a > 0$, and the remaining variables are defined as in the introduction. An attractive property of equation (4) is that it nests the benchmark quadratic loss function used in the basic Brainard model, collapsing to it for $a \rightarrow 0$. When $a > 0$, equation (4) is asymmetric, with a greater loss when financial stability is low ($x < x^*$) than when it is too high ($x > x^*$). Further, the greater a is, the greater this differential, and the more costly low financial stability is relative to high financial stability. Figure 1 plots the financial stability loss function for a range of values of a (holding $k = 0$). As can be seen, for larger values of a, the costs of very low financial stability (to the left of 0) can be significantly higher than the costs of very high financial stability (to the right of 0).

[3]We would like to thank David Aikman for suggesting this functional form.

As in the basic Brainard model, the policy tool k influences financial stability x in a linear way,

$$x = bk + u, \tag{5}$$

with uncertainty over the impact of the tool, b, and shocks to financial stability, u. For tractability we assume that $b \sim N\left(\bar{b}, \sigma_b^2\right)$, $u \sim N\left(0, \sigma_u^2\right)$, and $Cov\left(b, u\right) = 0$. Using equation (5) and the properties of the log-normal distribution, the objective function can be written as[4]

$$W = -\frac{1}{a^2}\left[exp\left(a\left(x^* - \bar{b}k\right) + \frac{a^2\left(\sigma_b^2 k^2 + \sigma_u^2\right)}{2}\right) - a\left(x^* - \bar{b}k\right) - 1\right] - \frac{\lambda}{2}k^2. \tag{6}$$

This function is strictly concave and has a unique global maximum $\tilde{k} > 0$ (see the appendix for the proof).

In the appendix we show that if the tool is expected to be sufficiently powerful, that is to say that $\bar{b}$ is large enough that a negative realization of b is a remote possibility, then

$$\frac{d\tilde{k}}{da} > 0. \tag{7}$$

That is, in the presence of uncertainty, greater asymmetry in the policymaker's financial stability objective leads to more active policy. When the losses from financial instability are sufficiently greater than those from too much stability and the policy tool is reasonably effective, the policymaker is more active to make sure low-stability outcomes are avoided. Such activism brings large benefits in reducing the cost of a financial crisis with relatively small costs if financial stability ends up being too high.[5]

[4]Specifically, we make use of the result that for $Z \sim N\left(\mu, \sigma^2\right) : \mathbb{E}\left[exp\left(Z\right)\right] = exp\left(\mu + \frac{1}{2}\sigma^2\right)$.

[5]In the working paper version of this paper, we also show that if $\bar{b}$ and a are sufficiently large, then $\frac{d\tilde{k}}{d\sigma_b^2} > 0$ and $\frac{d\tilde{k}}{d\sigma_u^2} > 0$, as the asymmetry in the objective prompts the policymaker to insure against downside risks by being more active.

2.2 Discussion

There is good reason to think macroprudential policymakers should have an asymmetric preference when it comes to financial stability. The substantial skew with regards to negative economic outcomes when financial stability risks materialize means that the costs of missed downside risks may be much larger than benefits of erring towards looser policy. Theoretical models that deliver endogenous crises via occasionally binding constraints (see, for example, Bianchi and Mendoza 2010, Jeanne and Korinek 2013, Korinek 2011, and Korinek and Simsek 2014) illustrate how severe asymmetries can emerge in stylized macroeconomic models and show that this can provide a normative justification for macroprudential policies that are contractionary in normal circumstances. The intuition for the findings above are similar: as the policymaker is much more averse to low financial stability outcomes, she buys insurance against low financial stability by tightening policy. The cost of this is an overly tight policy choice in situations when the financial system is stable.

We see this reasoning in the mandates of macroprudential policymakers. In the United Kingdom, for instance, the Financial Policy Committee's primary objective is to identify, monitor, and take action to remove systemic risks, thus indicating a focus on the downside of risks to the financial system. In contrast, the mandate for monetary policy tends to be to avoid both high and low inflation, and in many cases the target is explicitly symmetric (see, for example, the mandate of the Sveriges Riksbank). However, there are reasons why asymmetries may also matter in the context of monetary policy. For example, if the zero lower bound exacerbates deflationary shocks, the policymaker will have a stronger preference for avoiding below-target inflation. On the other hand, inflation bias tends to be positive, hence central banks could be concerned that above-target inflation will be more damaging to their credibility. Patton and Timmermann (2007) provide some empirical evidence that monetary policymakers have asymmetric preferences. Specifically they show that the Federal Reserve's revealed preferences, based on forecast errors, are more averse to lower realizations of output growth.

3. Targeting the Worst-Case Outcome (Robustly)

The conclusions from Brainard (1967), and indeed most modeling frameworks where policymakers maximize the expectation of an objective, rely on policymakers being able to assign probabilities to potential future scenarios. The previous section considered what happened if more weight was placed on negative realizations but the policymaker was able to assess the likelihood of such realizations occurring. However, this could be an unrealistic way of describing the uncertainty faced by macroprudential policymakers. It is not always possible to say with confidence how likely a relevant outcome will be. There is an inescapable need for policymakers to make judgments about the state of the world that cannot be backed by statistical analysis (Svensson 2002). This sort of unquantifiable uncertainty is sometimes referred to as fundamental or Knightian uncertainty. And it is a form of uncertainty that may be particularly troublesome for macroprudential policymakers given the innovation and increasing complexity of the financial system, which can make risks difficult to quantify. Furthermore, the objective of macroprudential policymakers is defined in terms of the resilience of the financial system, suggesting a focus upon rare events whose likelihood is undetermined.

A popular method among economists for incorporating fundamental uncertainty and concerns over severe outcomes into policymaking decisions is to rely upon a robust control approach. This approach favors policies that avoid large losses across scenarios regardless of how likely any given scenario is. In practice this is implemented by what is called a min-max framework (Wald 1950): a policymaker sets policy to minimize her losses assuming that the parts of the problem she is fundamentally uncertain about are chosen to maximize her loss. This is equivalent to making the worst-case scenario as palatable as possible.

3.1 Modeling Framework

To introduce a robust control motive into the Brainard model, we modify the relationship between k and x to be

$$x = bk + u + v, \tag{8}$$

where with the exception of v, all variables are defined as in equation (1) in the introduction. The term v captures what we refer to as fundamental uncertainty over x, the level of financial stability. This is a source of uncertainty which is ambiguous, and the policymaker is unable to attach a probability distribution to it. A robust approach to dealing with the uncertainty captured by v is to choose policy on the assumption that the worst outcome has happened so as to avoid large losses. To implement this, the min-max framework has the policymaker choosing k to maximize its objective conditional on "nature" choosing v to minimize the objective

$$max_k min_v - \frac{1}{2}\mathbb{E}[(x - x^*)^2 + \lambda k^2] + \frac{\theta}{2}v^2,\, \theta > 1. \tag{9}$$

Both nature and the policymaker choose their action conditional on each other's choice, and a pure-strategy Nash equilibrium solves the model. In effect, the policymaker and nature play a phantom game with each other, and nature behaves as a fictitious evil agent that chooses v to ensure the worst-case outcome occurs. An important element to the robust control problem is the term θv^2. Using the definitions of Hansen and Sargent (2008), the parameter θ is called entropy and, at a fundamental level, both represents how wrong the policymaker thinks they can be about the true level of x and penalizes extreme realizations of v.[6] From the perspective of the optimization problem, θv^2 simply serves as a constraint on nature when choosing how bad the worst-case scenario can be. The constraint $\theta > 1$ ensures that nature's problem is convex and so a finite v is chosen.

It is straightforward to show that the solution for nature, taking the actions of the policymaker as given, is given by

$$v = \frac{\bar{b}k - x^*}{\theta - 1}. \tag{10}$$

Thus as θ goes to infinity, nature is constrained to choose $v = 0$ and the policymaker ignores the fundamental uncertainty; conversely, as

[6] An alternative interpretation, as in Hansen and Sargent (2001), is that v represents model misspecification: the policymaker does not know the true model driving financial stability and thus v reflects the perturbation of the true model from the model that is being relied upon.

$\theta \to 1$, the worst-case scenario becomes increasingly bad and the policymaker becomes more and more concerned about the ambiguity over v.

The policymaker's optimal choice of k, taking the actions of nature as given, is given by

$$k = \frac{\bar{b}x^* - \bar{b}v}{\bar{b}^2 + {\sigma_b}^2 + \lambda}. \tag{11}$$

Solving for the Nash equilibrium gives the robust control solution to k as

$$k = \frac{\frac{\theta}{(\theta-1)}\bar{b}x^*}{\frac{\theta}{(\theta-1)}\bar{b}^2 + {\sigma_b}^2 + \lambda}. \tag{12}$$

Inspecting this solution, it is clear that as $\theta \to \infty$ and fundamental uncertainty disappears, the robust control solution for k is the same as the Brainard solution given by equation (3) in the introduction. Note also that the equilibrium realization for v is

$$v = \frac{-x^*({\sigma_b}^2 + \lambda)}{\theta\bar{b}^2 + (\theta - 1)({\sigma_b}^2 + \lambda)} < 0. \tag{13}$$

The negative value of v means that the worst-case scenario, in this model, is one where there is a negative shock to financial stability. This is intuitive in the context of a macroprudential policymaker.

The main result of this section is that policy becomes more activist when there is greater fundamental uncertainty,[7]

$$\frac{dk}{d\theta} < 0. \tag{14}$$

Thus, in particular, as the solution coincides with Brainard when $\theta \to \infty$ and fundamental uncertainty disappears, the policymaker is more activist than in the case of Brainard when there is fundamental uncertainty (i.e., for $\theta \in (1, \infty)$).

The simple explanation for this finding is that increasing fundamental uncertainty makes the worst-case outcome, which the policymaker is preparing for, worse. As the worst-case outcome is one

[7]Recall that fundamental uncertainty increases as θ decreases.

Figure 2. Graphical Illustration of the Robust Control Model

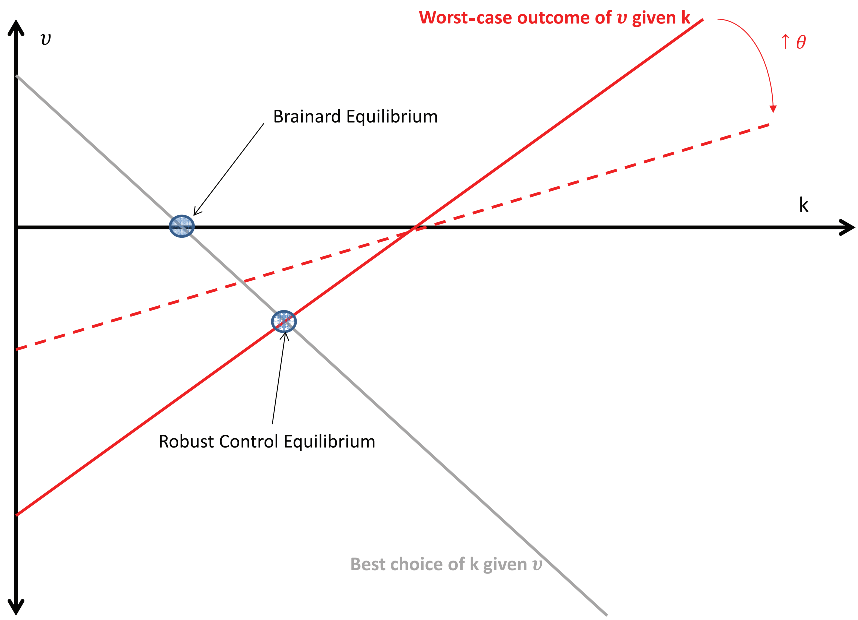

where there is a large negative shock to financial stability, this requires a tighter policy stance in response.

Figure 2 illustrates the intuition behind the robust control model graphically. The downward-sloping light grey line gives the best choice of k given v (given by equation (11)). The intuition for the downward slope is straightforward: a negative v implies a negative shock to financial stability and hence policy should be tightened to stabilize x. The upward-sloping line captures the "prepare for the worst" nature of the robust control problem: it is the value of v that gives the worst possible outcome to the policymaker given k and subject to the entropy constraint (given by equation (10)). This line is upward sloping because if loose policy is set, the worst-case outcome is a negative shock to financial stability. A robust policy choice is to set k where the two lines intersect: this means that the policymaker is choosing a policy which gives the best result if the very worst happens. In the standard Brainard problem in the introduction, $v = 0$, so the Brainard solution is equivalent to where the grey line crosses the k-axis. In the figure the intersection between the two lines is at a higher level of k than the Brainard point. Hence the robust control policy is more aggressive than Brainard would suggest. The intuition for this lies with the idea that more policy action is required if the most severe scenario—a negative financial stability shock—occurs, and the higher the fundamental uncertainty, the more severe is the worst case.[8] Increasing θ reduces fundamental uncertainty; this flattens the upward-sloping line[9] (shown by a dashed line in figure 2) and means that the worst-case outcome is not as bad, implying that policy does not need to be so activist.

3.2 Discussion

While a robust approach to policymaking usually suggests policy should be more active, it may say little about the direction in which policy should be more aggressive. To see this, consider the response of fiscal policy in the event of a severe recession. One extreme outcome may be high unemployment, negative growth, and deflation. Therefore, the robust response may be to embark upon aggressive fiscal expansion. However, an alternative scenario is that the recession puts sufficient strain on public finances to bring the government's solvency into question, leading to rising sovereign risk premia and additional stress on the financial system and the economy. So an alternative robust policy may be an aggressive tightening of the fiscal stance. Therefore, a policy designed to avoid either extreme outcome requires an aggressive shift in the policy stance, but the exact direction depends on which scenario policymakers are trying to avoid.

At face value, for the macroprudential context this lack of direction may seem less of an issue. Macroprudential policymakers are chiefly concerned by the extreme tail risks posed by financial crises.

---

[8] This is a modeling assumption. It is not the case that all robust control problems lead to a more aggressive policy, but it is a typical feature of this sort of model. See Onatski and Stock (2002) for a discussion in the context of monetary policy. Barlevy (2009) offers some counterexamples.

[9] Equation (10) can be rewritten as $k = \frac{v(\theta-1)+x^*}{b}$. Hence, nature's best response function intersects the k-axis at the point $k = \frac{x^*}{b}$ which is independent of θ. Therefore, any change in θ is a rotation of nature's best response line around the point $(\frac{x^*}{b}, 0)$, as illustrated in figure 2.

This asymmetry means that, at most points in time, a policymaker behaving in a robust way would be hawkish in setting her macroprudential tools. However, there is always a risk that under certain conditions severe negative outcomes can emerge if policy is set too tightly. In the spirit of Lucas (1987), one can imagine that the worst outcome of macroprudential policy is to lower the long-term trend growth rate. The question would then be whether an overly burdensome regulatory regime had a greater impact on long-term growth than financial crises.

Despite its popularity in the academic literature, a robust control approach can appear abstract when applied to practical policymaking. It has real-world applications nonetheless. For example, macroprudential policymakers often have a large number of potential indicators of risks to the financial system: a robust policymaking strategy would be to pay most attention to indicators that are signaling problems rather than focus on a measure of central tendency such as a simple average. Second, the principle of calibrating policy to prepare for severe outcomes is already embodied within many macroprudential frameworks via stress testing. Stress tests by definition provide a sense of economic outcomes if an extreme scenario emerges.

4. Learning about the Effects of Policy

A natural way to respond to uncertainty over a policy instrument is to attempt to learn about it. Research using evidence from other countries or from natural experiments, or by using calibrated theoretical models, can fill this gap. However, there is rarely a perfect substitute for using the instrument itself. Furthermore, in the macroprudential context, the framework for making and communicating changes to the tools is also new. Equivalent instruments have been used for microprudential purposes (such as bank capital requirements), but the signaling value (see section 5) and systemwide consequences of a macroprudential action may lead to a different impact, limiting what can be learned from previous uses of the tools.

In this section, we adapt the Brainard model to allow the policymaker to learn about the impact of her tool by observing what happens when she moves it. There is a trade-off: being active with a policy tool today leads to additional volatility, but by observing the tool's effect the policymaker will be less uncertain in future. This means that the policymaker has an incentive to be more active initially. While it may not be possible to directly measure or observe financial stability, so long as the policymaker can observe signals which convey some information about risks to the stability of the financial system (e.g., prices in financial markets), they will be able to learn about the effectiveness of their tool.

4.1 Modeling Framework

Here we lay out a two-period ($t = 1, 2$) version of the model presented in the introduction. Let

$$W_t := -\frac{1}{2}\mathbb{E}[(x_t - x^*)^2 + \lambda\,(k_t)^2]. \tag{15}$$

And suppose the policymaker's objective in period 1 is now to maximize

$$W_1 + \delta W_2 : \delta > 0. \tag{16}$$

As before, the policy instrument k_t has an uncertain impact on x_t described by

$$x_t = bk_t + u_t. \tag{17}$$

The parameter δ determines how much the policymaker values future periods. The other parameters share the interpretation they have in the introduction. For tractability we suppose that b and u_t are normally and independently distributed with the following initial priors: $b \sim N\left(\bar{b}, \sigma_b^2\right)$ and $u_t \sim N\left(0, \sigma_u^2\right)$.

The crucial assumption underpinning this two-period model is that by setting policy in period 1 the policymaker learns something about b in period 2. To embed this feature in the model, we have to specify the information available to the policymaker in each period. We assume that in period 1, when k_1 is chosen, the policymaker has the same information as in the Brainard model. At the end of period 1 the policymaker observes x_1, which, given k_1, provides an additional signal over b, reducing the uncertainty around the impact of the tool. However, since u_1 is not observed, the policymaker cannot perfectly distinguish between movements in x_1 that are due to k_1 or due to u_1. This results in a simple signal-extraction problem where the policymaker updates the expectation she has of b and uncertainty that surrounds that expectation (as shown in the appendix).

$$b|x_1 \sim N\left(\bar{b} + \frac{\sigma_b^2 k_1}{(k_1)^2\sigma_b^2 + \sigma_u^2}(x_1 - \bar{b}k_1),\ \frac{\sigma_b^2\sigma_u^2}{(k_1)^2\sigma_b^2 + \sigma_u^2}\right) \tag{18}$$

Equation (18) implies that the updated variance in period 2 is strictly less than in period 1 when $k_1 > 0$ ($Var(b|x_1) < \sigma_b^2$); thus uncertainty is reduced between the two periods. Moreover, increasing k_1 reduces $Var(b|x_1)$: the policymaker becomes more certain in the second period the more she acts in the first. The problem in period 2 is almost identical to the static Brainard model of the introduction, and the policymaker tries to maximize W_2 conditional upon beliefs about the mean and variance about b. The only difference is that now those beliefs are given by $\mathbb{E}(b|x_1)$ and $Var(b|x_1)$, which depend on the choice of k_1 in the previous period.

$$k_2 = \frac{\mathbb{E}(b|x_1)x^*}{\mathbb{E}(b|x_1)^2 + Var(b|x_1) + \lambda} \tag{19}$$

As a result, the policymaker faces a trade-off: actively using a tool in period 1 may not be immediately desirable due to the Brainard result from the introduction, but being active will make the policymaker better off in the future via lower uncertainty. To illustrate how this trade-off manifests and interacts with other parameters in the model, in figure 3 we present the optimal choice of k_1 across different levels of δ and ${\sigma_u}^2$.[10] These values of k_1 are presented in comparison to the optimal choice if the policymaker faced no uncertainty about b and if policy was set using the Brainard model.

Consider the left-hand panel of figure 3 first. We see the result that increasing δ monotonically increases k_1. The intuition for this is straightforward: as δ increases, the policymaker puts more weight on future outcomes. Hence, learning about how policy works in period 2 becomes a greater priority leading to a more activist policy. If $\delta = 0$

[10] It is clear from equations (18) and (19) that the policymaker's maximization problem can be set up in terms of k_1 only. This results in a highly non-linear function which is solved numerically using Monte Carlo methods.

Figure 3. Policy Choice with Learning

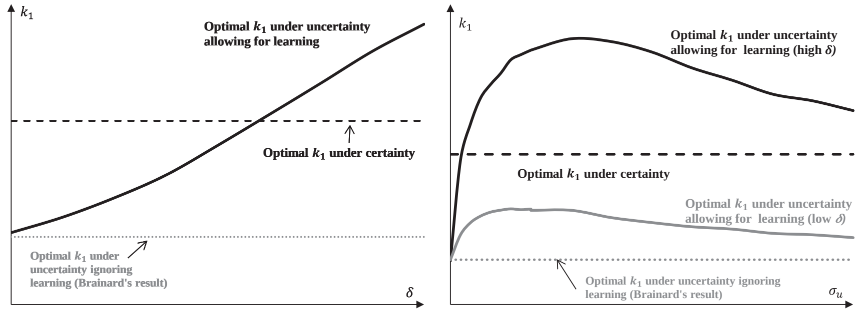

(i.e., there is no learning motive), the best choice for k_1 is simply in line with the Brainard model. For high values of δ (greater than 1) we can see that k_1 can be greater than it would be under certainty about b.[11] It is also possible to show, as we do in the appendix, that if the policymaker only cares about the second period (as they would when $\delta \to \infty$), then it is optimal for $k_1 \to \infty$ when there is no motive to stabilize k.

The right-hand panel of figure 3 considers how policy varies with ${\sigma_u}^2$. The irrelevance of uncertainty over the state of the economy for the policy choice does not extend to the learning case; instead we see a hump-shaped pattern with respect to ${\sigma_u}^2$. This is because how much can be learned about b by altering k_1 has a non-linear relationship with the uncertainty over u_1. If ${\sigma_u}^2$ is zero, there is nothing more that can be learned from altering k_1, as b is observed perfectly at the end of period 1 and thus the optimal k_1 is the same as Brainard. As ${\sigma_u}^2$ rises from zero, the uncertainty over b in period 2 increases, hence the policymaker chooses to be more active to learn more about the coefficient. However, as ${\sigma_u}^2$ becomes increasingly large, the ability of x_1 to be informative about b diminishes and the policymaker becomes less able to learn and is therefore less active.

[11]In reality policymakers will need to make decisions over many periods, thus the loss at period 2 could be thought of as a reduced-form way of capturing the continuation value of the problem. This would suggest that a large δ, potentially greater than 1, is appropriate.

At the limit ${\sigma_u}^2 \to \infty$, there is nothing that can be learned from x_1 and policy reverts to the Brainard solution.[12]

4.2 Discussion

The desire to learn is intimately related to how uncertain the policymaker thinks they will be in future. From today's perspective the macroprudential toolkit is largely untested. This is in contrast to the current state of monetary policy which can draw upon decades of academic research into its effects, substantial policy experience, and data sets that contain the necessary variation. In the macroprudential context, the extant empirical literature relies on difference-in-differences based empirical strategies, which can assess the impact of a small number of quasi-natural experiments but cannot account for differing macroeconomic climates or general equilibrium consequences. However, in the future, research into macroprudential policy will eventually come on stream and contribute to developing a better understanding of how the tools operate. In the meantime, policymakers today have to decide on the extent to which the current state of the literature can substitute for seeing the tool in action. Furthermore, using a policy generates both operational experience and the variation required to conduct detailed empirical studies.

The nature of the financial cycle means that the ability of policymakers to learn about the effectiveness of macroprudential policy differs substantially from the case of monetary policy. On the one hand, the relatively infrequent occurrence of financial crises means that, over time, there are relatively few events to learn from. On the other hand, the length of financial cycles means that, if policymakers believe that they can readily observe financial stability, there could be more opportunity to learn about the impact of policy tools ahead of crises.

The immediate practical issue with using a tool to learn is one of communication and political sensitivity. A macroprudential policymaker would probably struggle to articulate to banks that it was forcing them to raise more capital in order to determine the economic consequences. Deliberately experimenting with the economy with no other goal in mind is highly inadvisable and should run counter to

[12] These results for σ_u^2 are shown formally in the appendix.

macroprudential policy objectives. However, it is not necessary to implement policy that is harmful in order to learn. Though policy should not be set with wild abandon, the learning mechanism provides an argument to suggest that there are benefits from not being too cautious.

5. Policy Transparency and Private-Sector Uncertainty

It is not only policymakers that face uncertainty with regards to macroprudential policymaking. The public is also uncertain about risks to financial stability, the impact of macroprudential policy, and how the policymaker herself will behave. Furthermore, the policymaker may have informational advantages over the public through access to confidential regulatory data, supervisory intelligence, and the results of stress tests. Thus a macroprudential policymaker needs to consider how her policy actions interact with the information available to the public.

In the models used in the introduction and the previous three sections, the private sector has responded to policy actions in a mechanical fashion. In this section, we allow for strategic interaction between the policymaker and the private sector. The setup moves beyond a simple stabilization problem into a framework where (private) risk-taking is an endogenous choice (albeit modeled in a stylized fashion). The transmission mechanism of policy is similar, as is the eventual objective of the policymaker, to that discussed in the previous sections. However, this framework allows us to explore how private-sector uncertainty affects the transmission mechanism and, hence, how policymakers should behave in response.

We assume that the policymaker is better informed about the stability of the financial system than the private sector. As such, the policy action is an additional signal for the private sector. Crucially, this signal is imperfect: we also assume that the private sector is not fully aware of the objectives of the policymaker and hence cannot perfectly distinguish actions caused by risks to financial stability from those caused by an aversion to crises. An example of this would be if banks did not know if the policymaker set a high countercyclical capital buffer rate because it was aware of a specific financial stability risk or simply had a high preference for avoiding a crisis.

5.1 Modeling Framework

The model takes the form of a Stackelberg game. The policymaker first chooses the level of her policy instrument, k; then, having observed k, the private sector chooses their desired level of risk-taking, x. The policymaker internalizes the impact of her policy choice on the private sector's response when choosing her desired level of k.

We assume the private sector chooses x in order to maximize the following objective:

$$max_x \mathbb{E}\left[\left(-\frac{1}{2}(x - x^+ + u)^2 + \alpha x k\right) |k\right], \tag{20}$$

where α is a parameter whose interpretation we will return to shortly. There are two stochastic variables in this setup: u and c. The former enters the private sector's loss function; the latter is not directly relevant for the private sector but affects the preferences of the policymaker. Both u and c are uncorrelated random variables with positive variance, and the private sector has priors $u \sim N(0, \sigma_u^2)$ and $c \sim N(0, \sigma_c^2)$. This functional form has similarities with the stabilization problem described in the previous sections, with three main differences. First, x is now a choice variable of the private sector and there is a positive private target for risk-taking denoted $x^+ > 0$, capturing, for example, instability bias due to limited liability. Second, for given k, u interferes with the optimal level x that the private sector would prefer to target. A higher value of u implies that lower x is desired. This is consistent with the idea that u is an orthogonal shock that negatively affects financial stability and would make the private sector want to take less risk. The third difference is the cross-term $\alpha x k$. This term is missing from the models considered in the previous sections, and its presence allows for strategic interaction between the private sector and the policymaker. We assume $\alpha < 0$, which implies that, at the margin, tighter policy encourages less risk-taking by the private sector (this would be the case if, for instance, tighter capital requirements increase the private sector's skin in the game).[13]

[13]Our main result can also go through with α positive, which could occur for instance if tighter policy reduces the severity of the left tail of economic outcomes, encouraging risk-taking. In that case, equilibrium existence requires the assumption that $2\alpha\beta < 1$, where β is defined below.

The private sector takes k as given and observes its value before making its decision. The policymaker cares about u but has superior information about its realization; hence, k serves as a signal to the private sector about u. Formally, we can write the private sector's first-order condition as

$$x = \alpha k + x^{+} - \mathbb{E}(u|k), \tag{21}$$

and so the private sector takes less risk as $\mathbb{E}(u|k)$ increases. The model is solved via the method of undetermined coefficients, and we postulate (and then verify) that the private sector's belief about u, having observed k, is linear in the choice of policy tool

$$\mathbb{E}(u|k) = \gamma_0 + \gamma_1 k. \tag{22}$$

The policymaker's problem is

$$max_k \left[-\frac{1}{2}(k - u - c)^2 + \beta x(k) k \right], \tag{23}$$

where from above, $x(k) = (\alpha - \gamma_1)k - \gamma_0 + x^{+}$. We also assume that the policymaker has full information (i.e., u and c are known). Again, we can think of a positive realization of u as a negative shock to financial stability which warrants tighter policy. In addition to u, the policymaker also has her own preferred target level for k, denoted c, which does not directly influence (and is unknown by) the private sector. Furthermore, we have a second cross-term: βxk. We assume that $\beta > 0$ to capture the idea that, at the margin, more risk-taking by the private sector makes tighter policy more desirable. Last, note that the policymaker internalizes that her decision over k will affect x, hence x enters the optimization problem as a function of policy.

These assumptions over preferences allow us to have a short exposition of the main mechanism we wish to highlight. The working paper version of this paper contains a larger model where preferences are more general (although still linear quadratic) and generates the same main result. The policymaker's problem has a unique maximum whenever

$$-\frac{1}{2} + \beta(\alpha - \gamma_1) < 0. \tag{24}$$

In the appendix we show that this condition is satisfied in equilibrium under the given parameter assumptions. Moreover, in equilibrium, $\gamma_1 > 0$. On this basis, taking the policymaker's first-order condition and solving for k gives

$$k = \frac{u + c + \beta(x^+ - \gamma_0)}{(1 - 2\beta(\alpha - \gamma_1))}. \tag{25}$$

From the perspective of the private sector, k is then normally distributed, and we can then solve for $\mathbb{E}(u|k)$ using the standard signal-extraction formula for normal distributions:

$$\mathbb{E}(u|k) = \mathbb{E}(u) + \frac{Cov(u,k)}{Var(k)}(k - \mathbb{E}(k)). \tag{26}$$

In the appendix we show that this is of the same form as equation (22), verifying the form of the postulated solution. We further show through equating coefficients that $\gamma_0 = \frac{-s\beta x^+}{1-\beta s}$ and

$$\gamma_1 = \frac{(1 - 2\beta\alpha)\, s}{1 - 2\beta s} > 0, \tag{27}$$

where $s = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}$ is a function of the signal-to-noise ratio, and equation (27) implies that we must also have $2\beta s < 1$. The private sector cares about u, but c, the policymaker's personal preference for tighter policy, does not influence its behavior. Hence, k is a useful signal to the extent that it provides information about u, but uncertainty about c adds noise to this signal. In particular, the greater the public-sector transparency about c, the lower σ_c^2 will be and the greater will be s. The greater the signal-to-noise ratio, the more informative k is as a signal about u. This is apparent upon inspecting equation (27), where we can see that $\frac{d\gamma_1}{ds} > 0$, consistent with the intuition that the sensitivity of $\mathbb{E}(u|k)$ to k is increasing in the signal-to-noise ratio.

In the appendix we also show that the policymaker's optimal choice is given by

$$k = \frac{1 - 2\beta s}{1 - 2\beta\alpha}\left(u + c + \frac{\beta x^+}{1 - \beta s}\right). \tag{28}$$

The main result of the model is then that the greater the signal-to-noise ratio, and hence the lower the private-sector uncertainty about policy preferences, the less active the policymaker is in responding to u. This is clear upon inspecting equation (28). The mechanism at work runs through the fact that the sensitivity of x to k is given by $\alpha - \gamma_1$. Recall that we assume that α is negative, meaning that absent any signaling channel, the private sector will take less risk as k is increased. The signaling channel amplifies the response of private-sector risk-taking to policy. When signaling is present, as s increases k becomes a more informative signal about the realization of u and consequently more weight is placed on the signal, resulting in γ_1 being larger. This in turn increases the responsiveness of x to k as the private sector view a larger part of the rise in k being due to a less stable financial system. Consequently, with the private sector taking less risk in response to a given policy choice, k does not need to be set as high, resulting in a lower sensitivity of k to u.

5.2 Discussion

This result highlights the importance of a policymaker ensuring that private agents understand why a policy action is taken, in order to maximize the signaling value of a policy decision. In practice, a macroprudential policymaker has a variety of ways of communicating to the private sector beyond her actions. The explanation of macroprudential policy decisions given alongside the action can help the public learn and extract an informative signal about the policymaker's views on financial stability risks. Alternatively, the macroprudential policymaker may signal their likely future actions in response to a buildup in risks by clarifying their "reaction function." The publication of core financial stability indicators may provide some clarity to the policymaker's reaction function. Communicating likely future actions may affect expectations about the evolution of the economy and hence influence behavior today.

Changes in private-sector uncertainty have affected the transmission mechanism of other policy tools. As noted by Beechey (2008), the news on May 6, 1997 that a Monetary Policy Committee would be created and the Bank of England would be granted operational independence with an explicit inflation target was followed by a substantial reduction in market-implied short- and long-term inflation expectations, arguably because this reduced uncertainty about the objectives of monetary policy. The forward rates of inflation compensation five to ten years ahead fell around 35 basis points on the day of the announcement, with a further decline over the following days. This in turn affected the transmission mechanism of monetary policy.

The obvious difference, however, between macroprudential and monetary policymakers is that in the context of the latter a precise, measurable target is available. Furthermore, the private sector has had more chance to learn about how monetary policy is conducted in practice and, correspondingly, policymakers' reaction functions.

6. Conclusion

As this paper has highlighted, there are several types of uncertainty and multiple channels through which uncertainty can affect policymaking. The well-known result of Brainard (1967) that policy should be more cautious in the presence of uncertainty does not hold in general, particularly for specific examples that are relevant to macroprudential policymakers. If anything, the need to learn about these relatively untested tools and the focus on avoiding tail risks speak to more active policymaking. Moreover, private-sector uncertainty over the financial stability objectives, preferences, and reaction function may diminish the potency of the signaling impact of macroprudential policy, requiring more active policy or communication. One limitation of the analysis here is that it takes place in context of static or two-period models. Hence it is a challenge to disentangle tighter steady-state macroprudential policy in response to increased risks from a more active response to economic shocks as they emerge. This would be an interesting area for future analysis.

Appendix 1. Asymmetric Objective Function

PROPOSITION 1. *The optimal choice of policy under asymmetry is given by the* $\tilde{k}$ *that satisfies*

$$\frac{-1}{a}\left(exp\left\{ax^* - a\bar{b}\tilde{k} + \tfrac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right\}\left(-\bar{b} + a\tilde{k}\sigma_b^2\right) + \bar{b}\right) - \lambda\tilde{k} = 0. \tag{29}$$

Further, $\tilde{k}\epsilon\left(0, \frac{\bar{b}}{a\sigma_b^2}\right)$.

Proof. The policy objective under asymmetry is W, as given by equation (6). The first derivative of W w.r.t. k is given by

$$-\left\{\frac{1}{a}\left[exp\left\{ax^* - a\bar{b}k + \tfrac{a^2}{2}\left(k^2\sigma_b^2 + \sigma_u^2\right)\right\}\left(-\bar{b} + ak\sigma_b^2\right) + \bar{b}\right] + \lambda k\right\}.$$

The second derivative of W w.r.t. k is given by

$$-\left\{\frac{1}{a^2}\left[e^{\left(ax^* - a\bar{b}k + \frac{a^2}{2}\left(k^2\sigma_b^2 + \sigma_u^2\right)\right)}\left\{-a\bar{b} + a^2k\sigma_b^2\right\}^2 \right.\right.$$
$$\left.\left. + e^{\left(ax^* - a\bar{b}k + \frac{a^2}{2}\left(k^2\sigma_b^2 + \sigma_u^2\right)\right)} a^2\sigma_b^2\right] + \lambda\right\}.$$

As this is always negative, W is a strictly concave function and hence equating the first-order condition (FOC) to 0 gives the unique global maximum.

Turning to the bounds on $\tilde{k}$, the first derivative, evaluated at $k = 0$, is

$$\frac{\bar{b}}{a}\left[exp\left(ax^* + \tfrac{a^2\sigma_u^2}{2}\right) - 1\right] > 0,$$

as all parameters are positive. Hence, as W is strictly concave, the derivative is positive iff $k < \tilde{k}$ and so we must have $0 < \tilde{k}$. Turning to the upper bound, if $k \geq \frac{\bar{b}}{a\sigma_b^2}$ then $-\bar{b} + ak\sigma_b^2 \geq 0$ and $\frac{dW}{dk} < 0$. However, as W is strictly concave, the derivative is negative iff $k > \tilde{k}$ and so we must have $\tilde{k} < \frac{\bar{b}}{a\sigma_b^2}$.

LEMMA 1. *Let $\tilde{k}$ be the policymaker's solution. Then $\frac{d\tilde{k}}{da} > 0$ iff the following expression is positive:*

$$e^{\left(ax^{*}-a\bar{b}\tilde{k}+\frac{a^{2}\left(\tilde{k}^{2}\sigma_{b}^{2}+\sigma_{u}^{2}\right)}{2}\right)}\left\{\frac{\bar{b}}{a}-\tilde{k}\sigma_{b}^{2}\right\}\left[\left(x^{*}-\bar{b}\tilde{k}\right)+a\left(\tilde{k}^{2}\sigma_{b}^{2}+\sigma_{u}^{2}\right)\right]$$
$$+\frac{\bar{b}}{a^{2}}\left[1-e^{\left(ax^{*}-a\bar{b}\tilde{k}+\frac{a^{2}\left(\tilde{k}^{2}\sigma_{b}^{2}+\sigma_{u}^{2}\right)}{2}\right)}\right]. \tag{30}$$

Proof. The optimal policy choice $\tilde{k}$ is defined implicitly by equation (29). Let $g(\tilde{k}\left(a\right),a)$ be the left-hand side, then $g(\tilde{k}\left(a\right),a)\equiv 0$. By the implicit function theorem, $\frac{d\tilde{k}}{da}=\frac{\partial g(\tilde{k}(a),a)}{\partial a}/-\frac{\partial g(\tilde{k}(a),a)}{\partial \tilde{k}}$. As the objective function is strictly concave, $\frac{\partial g(\tilde{k}(a),a)}{\partial \tilde{k}}<0$, as $g()$ is the first derivative of the objective function. Thus $\frac{d\tilde{k}}{da}>0$ iff $\frac{\partial g(\tilde{k}(a),a)}{\partial a}>0$. Taking the derivative of $g()$ w.r.t. a gives the result.

We first establish a useful lemma before turning to the main proposition.

LEMMA 2. *The following relationships hold as the expected effectiveness of the policy tool, $\bar{b}$, becomes arbitrarily large:*

$$\lim_{\bar{b}\to\infty}\tilde{k}=0 \tag{31}$$

$$\lim_{\bar{b}\to\infty}\frac{\bar{b}}{a^{2}}\left[1-exp\left(ax^{*}-a\bar{b}\tilde{k}+\tfrac{a^{2}}{2}\left(\tilde{k}^{2}\sigma_{b}^{2}+\sigma_{u}^{2}\right)\right)\right]=0 \tag{32}$$

$$\lim_{\bar{b}\to\infty}\bar{b}\tilde{k}=x^{*}+\frac{a\sigma_{u}^{2}}{2}. \tag{33}$$

Proof. To prove result (31), we establish that $\bar{b}\tilde{k}$ is bounded. We have established that $\bar{b}\tilde{k}>0$, so it is sufficient to show that $\exists S\geq 0:\bar{b}\tilde{k}\leq S\ \ \forall\bar{b}>0$. If S exists, it follows that result (31) holds.[14]

[14]To see this, note that if $\exists S\in\mathbb{R}_{+}:0<\bar{b}\tilde{k}\leq S\ \forall\bar{b}$, then $0<\tilde{k}\leq\frac{S}{\bar{b}}$ and hence as $\lim_{\bar{b}\to\infty}\frac{S}{\bar{b}}=0$, by the sandwich theorem we have $\lim_{\bar{b}\to\infty}\tilde{k}=0$.

To show that S exists, we suppose for a contradiction that it does not. Then $\bar{b}\tilde{k} \to \infty$ as $\bar{b} \to \infty$. We rewrite FOC (29) as

$$\frac{-1}{a}\left(e^{\left(ax^*-a\bar{b}\tilde{k}+\frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2+\sigma_u^2\right)\right)}a\tilde{k}\sigma_b^2 + \bar{b}\left\{1-e^{\left(ax^*-a\bar{b}\tilde{k}+\frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2+\sigma_u^2\right)\right)}\right\}\right) - \lambda\tilde{k} = 0. \tag{34}$$

To construct the contradiction, we first establish, under the supposition $\bar{b}\tilde{k} \to \infty$ as $\bar{b} \to \infty$, that

$$exp\left(ax^* - a\bar{b}\tilde{k} + \tfrac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right) \to 0 \text{ as } \bar{b} \to \infty. \tag{35}$$

To do this, we make use of the inequality $\tilde{k} \leq \frac{\bar{b}}{a\sigma_b^2}$, and thus the exponent in equation (35) must adhere to the following inequality:

$$\begin{aligned} ax^* - a\bar{b}\tilde{k} + \tfrac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right) &\leq ax^* + \frac{a^2\sigma_u^2}{2} - a\bar{b}\tilde{k} + \frac{a^2\sigma_b^2}{2}\tilde{k}\left(\frac{\bar{b}}{a\sigma_b^2}\right) \\ &= ax^* + \frac{a^2\sigma_u^2}{2} - \frac{a}{2}\bar{b}\tilde{k}. \end{aligned} \tag{36}$$

As $\bar{b}\tilde{k}$ is unbounded by assumption, the right-hand side of equation (36) tends to $-\infty$ as $\bar{b} \to \infty$ and hence so too does the left-hand side. Thus, under our supposition, equation (35) holds.

Now, equation (35) implies $\bar{b}\left\{1-exp\left(ax^*-a\bar{b}\tilde{k}+\frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2+\sigma_u^2\right)\right)\right\} \to \infty$.

Further, $exp\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right) a\tilde{k}\sigma_b^2 > 0$ and $\tilde{k} > 0$. Thus, the left-hand side of equation (34) tends to $-\infty$ as $\bar{b} \to \infty$, a contradiction, as the FOC no longer holds. Hence, we must have $\bar{b}\tilde{k}$ bounded, and thus result (31) holds.

Turning to the proof of result (32), rearranging equation (34), we have

$$\frac{\bar{b}}{a^2}\left[1 - e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)}\right]$$
$$= -\frac{\tilde{k}}{a}\left(\lambda + \sigma_b^2 e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)}\right).$$

The right-hand side $\rightarrow 0$ as $\bar{b} \rightarrow \infty$ because $\tilde{k} \rightarrow 0$ whilst, with $\tilde{k}$, $\bar{b}\tilde{k}$ bounded, the exponential term is also bounded.

Thus, we must have $\underset{\bar{b}\rightarrow\infty}{lim}\frac{\bar{b}}{a^2}\left[1 - e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)}\right] = 0$, showing result (32).

Finally, we must also then have $\underset{\bar{b}\rightarrow\infty}{lim}\, ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right) = 0$. Given $\underset{\bar{b}\rightarrow\infty}{lim}\,\tilde{k} = 0$, it follows that $\underset{\bar{b}\rightarrow\infty}{lim}\,\bar{b}\tilde{k} = x^* + \frac{a\sigma_u^2}{2}$. This proves result (33).

Using the lemma, we prove the main proposition of this section.

PROPOSITION 2. *When $\bar{b}$ is sufficiently large, policy becomes more active under greater asymmetry:* $\frac{d\tilde{k}}{da} > 0$.

Proof. To prove the proposition, we establish that equation (30) is positive as $\bar{b} \rightarrow \infty$. From the prior lemma,

$$\underset{\bar{b}\rightarrow\infty}{lim}\frac{\bar{b}}{a^2}\left[1 - e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)}\right] = 0$$
$$\underset{\bar{b}\rightarrow\infty}{lim}\, e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)} = 1.$$

Moreover,

$$\lim_{\bar{b}\to\infty} \left(x^* - \bar{b}\tilde{k}\right) + a\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right) = \left(x^* - \left(x^* + \frac{a\sigma_u^2}{2}\right)\right) + a\sigma_u^2$$
$$= \frac{a\sigma_u^2}{2} > 0.$$

As $\lim_{\bar{b}\to\infty} \tilde{k} = 0$, it follows that $\frac{\bar{b}}{a} - \tilde{k}\sigma_b^2 \to \infty$ and hence (30)$\to \infty$ as $\bar{b} \to \infty$. Thus, for sufficiently large $\bar{b}$, $\frac{d\tilde{k}}{da} > 0$.

Appendix 2. Learning about the Effects of Policy

LEMMA 3. *For a given choice of k_1, the posterior mean and variance of b (having observed x_1) are given by*

$$b|x_1 \sim N\left(\bar{b} + \frac{{\sigma_b}^2 k_1}{(k_1)^2{\sigma_b}^2 + {\sigma_u}^2}(x_1 - \bar{b}k_1),\ \frac{{\sigma_b}^2{\sigma_u}^2}{(k_1)^2{\sigma_b}^2 + {\sigma_u}^2}\right). \quad (37)$$

Proof. Given the distributional assumptions on b, u and given that k_1 is non-stochastic as chosen by the policymaker, we have that $x_1 \sim N\left(\bar{b}k_1, (k_1)^2\sigma_b^2 + \sigma_u^2\right)$. Then, as b, x_1 are both normally distributed, we have that

$$b|x_1 \sim N\left(\bar{b} + \frac{Cov\left(b, x_1\right)}{Var\left(x_1\right)}\left(x_1 - \mathbb{E}\left(x_1\right)\right),\ \sigma_b^2 - \frac{\left(Cov\left(b, x_1\right)\right)^2}{Var\left(x_1\right)}\right). \quad (38)$$

Now, $Cov\left(b, x_1\right) = Cov\left(b, bk_1 + u\right) = k_1\sigma_b^2$; applying this to equation (38) and rearranging completes the proof.

LEMMA 4. *Suppose $\sigma_u^2 = 0$ or $\sigma_u^2 \to \infty$; then k_2 is independent of k_1 and the solution for k_1 is as in the static case.*

Proof. From equation (37), it is clear that when $\sigma_u^2 \to \infty$, $b|x_1 \sim N\left(\bar{b}, {\sigma_b}^2\right)$ and so k_2 is independent of k_1 and hence the solution for k_1 is as in the static Brainard case.

Now suppose $\sigma_u^2 = 0$. Then from equation (37) $b|x_1$ has zero variance, and so in the second period the value of b is known perfectly for any choice of $k_1 \neq 0$. Thus, both $b|x_1$ and $b^2|x_1$ are known perfectly in the second period. Hence k_2 is independent of k_1 and thus the optimal solution for k_1 will be as in the static Brainard case.

PROPOSITION 3. *When $\lambda = 0$, when the policymaker only cares about the second period's welfare, it is optimal to set k_1 arbitrarily large.*

Proof. The relevant objective of the policymaker is now $\mathbb{E}(W_2)$, where $\mathbb{E}(.)$ denotes the prior expectation of the policymaker before x_1 is realized. $\mathbb{E}(W_2)$ has no closed-form solution; however, using the law of iterated expectations we can write $\mathbb{E}(W_2) = \mathbb{E}(\mathbb{E}(W_2 \mid x_1))$, where $\mathbb{E}(. \mid x_1)$ is the posterior expectation of the policymaker once x_1 is realized. As shown below, $E(W_2 \mid x_1)$ can be written in closed form. Therefore, the logic of the proof will be that, for $\lambda = 0$, $\mathbb{E}(W_2 \mid x_1)$ is maximized when $k_1 \to \infty$ for *any* realization x_1. In turn, this implies that $\mathbb{E}(W_2)$ is maximized when $k_1 \to \infty$.

Once x_1 is realized, we know that the optimal choice of k_2 is given by

$$k_1 = \frac{\mathbb{E}(b \mid x_1)x^*}{\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda}.$$

Substituting this into $\mathbb{E}(W_2 \mid x_1)$ yields

$$\mathbb{E}(W_2 \mid x_1) = -\frac{1}{2}E\left[\left(\frac{b\mathbb{E}(b \mid x_1)x^*}{\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda} + u_2 - x^*\right)^2 \right.$$
$$\left. + \frac{\lambda\mathbb{E}(b \mid x_1)^2(x^*)^2}{(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda)^2} \mid x_1\right].$$

Expanding the brackets gives $\mathbb{E}(W_2 \mid x_1) =$

$$-\frac{\mathbb{E}(b^2 \mid x_1)\mathbb{E}(b \mid x_1)^2(x^*)^2}{2(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda)^2} + \frac{\mathbb{E}(b \mid x_1)^2(x^*)^2}{\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda}$$
$$-\frac{\lambda\mathbb{E}(b \mid x_1)^2(x^*)^2}{2(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda)^2} + t.i.p.,$$

where $t.i.p.$ denotes terms independent of policy. We can then factorize to get

$$\mathbb{E}(W_2 \mid x_1) = -\frac{\mathbb{E}(b \mid x_1)^2 (x^*)^2}{2(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda)^2} \big(\mathbb{E}(b^2 \mid x_1) \\ - 2(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda) + \lambda\big) + t.i.p.$$

Using the fact that $\mathbb{E}(b^2 \mid x_1) = \mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1)$, the expression simplifies to

$$\mathbb{E}(W_2 \mid x_1) = \frac{\mathbb{E}(b \mid x_1)^2 (x^*)^2}{2(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda)} + t.i.p.$$
$$= \frac{(x^*)^2}{2\left(1 + \frac{Var(b|x_1)+\lambda}{\mathbb{E}(b|x_1)^2}\right)} + t.i.p.$$

Now when $\lambda = 0$, the maximum value that $\mathbb{E}(W_2 \mid x_1)$ takes is $\frac{(x^*)^2}{2}$, as all terms are positive. This occurs when $Var(b \mid x_1) = 0$ for any given realization of x_1 (except the knife-edge case where $\mathbb{E}(b \mid x_1)^2 = 0$ and $\mathbb{E}(W_2 \mid x_1)$ is independent of the policy choice). As described in the main text, $Var(b \mid x_1) = \sigma_b^2 \sigma_u^2 / ((k_1)^2 \sigma_b^2 + \sigma_u^2)$; therefore $Var(b \mid x_1) \to 0$ as $k_1 \to \infty$. Hence, $\mathbb{E}(W_2 \mid x_1)$ tends to its maximum when $k_1 \to \infty$. By the logic above, this implies that $\mathbb{E}(W_2)$ also attains its maximum. This completes the proof.

Appendix 3. Policy Transparency and Private-Sector Uncertainty

PROPOSITION 4. *Given our assumptions over parameters, a solution exists where private-sector beliefs are of the form* $E(u|k) = \gamma_0 + \gamma_1 k$. *Moreover,* $\gamma_1 > 0$ *and* $\frac{dk}{du}$ *is decreasing in* $s = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}$. *Further, the optimal policy choice under private-sector uncertainty is given by*

$$k = \frac{1 - 2\beta s}{1 - 2\beta\alpha} \left(u + c + \frac{\beta x^+}{1 - \beta s} \right).$$

Proof. Suppose that the private sector has a posterior expectation of the form $\mathbb{E}(u|k) = \gamma_0 + \gamma_1 k$. As shown in the text, so long as

$-\frac{1}{2} + \beta(\alpha - \gamma_1) < 0$, the optimality condition of the policymaker is then given by

$$k = \frac{u + c - \beta(\gamma_0 - x^+)}{(1 - 2\beta(\alpha - \gamma_1))}.$$

As u, c are normally distributed, so too is k and so, using the formula for conditional normal distributions,

$$\mathbb{E}(u|k) = \mathbb{E}(u) + \frac{cov(u,k)}{Var(k)}(k - \mathbb{E}(k)).$$

Now, with u, c uncorrelated,

$$\frac{cov(u,k)}{Var(k)} = \frac{\frac{\sigma_u^2}{(1-2\beta(\alpha-\gamma_1))}}{\frac{\sigma_u^2+\sigma_c^2}{(1-2\beta(\alpha-\gamma_1))^2}} = (1 - 2\beta(\alpha - \gamma_1))\frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}$$

$$\mathbb{E}(k) = \frac{-\beta(\gamma_0 - x^+)}{(1 - 2\beta(\alpha - \gamma_1))}.$$

Thus

$$\begin{aligned}\mathbb{E}(u|k) &= (1 - 2\beta(\alpha - \gamma_1))\frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}\left(k + \frac{\beta(\gamma_0 - x^+)}{(1 - 2\beta(\alpha - \gamma_1))}\right) \\ &= \frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}\left((1 - 2\beta(\alpha - \gamma_1))k + \beta(\gamma_0 - x^+)\right).\end{aligned}$$

This confirms that the postulated linear form for the private sector's posterior belief, $\mathbb{E}(u|k) = \gamma_0 + \gamma_1 k$, is correct. Equating the coefficients on k gives

$$\gamma_1 = (1 - 2\beta(\alpha - \gamma_1))\frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}.$$

Define $s = \frac{\sigma_u^2}{\sigma_u^2+\sigma_c^2}$, where s is a function of the signal-to-noise ratio, and $s \in (0, 1)$. Rearranging this equation gives

$$\gamma_1 (1 - 2\beta s) = (1 - 2\beta\alpha)\, s.$$

As $2\beta\alpha < 0 < 1$, the right-hand side is non-zero, and thus we must have $\gamma_1 \neq 0$, $(1-2\beta s) \neq 0$, giving

$$\gamma_1 = \frac{(1-2\beta\alpha)\, s}{1-2\beta s}.$$

Now,

$$\begin{aligned} 1-2\beta(\alpha-\gamma_1) &= \frac{(1-2\beta s) + 2\beta\,(1-2\beta\alpha)\, s - 2\beta\alpha\,(1-2\beta s)}{1-2\beta s} \\ &= \frac{(1-2\beta\alpha)\,[(1-2\beta s)+2\beta s]}{1-2\beta s}. \end{aligned}$$

Tidying gives

$$1-2\beta(\alpha-\gamma_1) = \frac{(1-2\beta\alpha)}{1-2\beta s}.$$

From equation (24), we have a valid equilibrium iff $1-2\beta(\alpha-\gamma_1) > 0$, and hence iff $\frac{(1-2\beta\alpha)}{1-2\beta s} > 0$. Thus, in any equilibrium in which private-sector conditional beliefs are of the linear form, we must have $\gamma_1 > 0$. Moreover, given this, with $2\alpha\beta < 0 < 1$ we must also have $2\beta s < 1$.

Now, the equilibrium policy choice will be

$$k = \frac{(u + c + \beta(x^+ - \gamma_0))\,(1-2\beta s)}{(1-2\beta\alpha)},$$

and so

$$\frac{dk}{du} = \frac{(1-2\beta s)}{(1-2\beta\alpha)} > 0,$$

which is decreasing in s. Finally, consider the coefficient γ_0. Equating the constant coefficient for $\mathbb{E}(u|k)$ gives $\gamma_0 = \beta s(\gamma_0 - x^+)$. Given that $2\beta s < 1$, $\beta s \neq 1$, and we must have $\gamma_0 = \frac{-s\beta x^+}{1-\beta s}$, and $x^+ - \gamma_0 = \frac{x^+}{1-s\beta}$.

References

Angeletos, G.-M., L. Iovino, and J. La'O. 2011. "Optimal Monetary Policy with Informational Frictions." NBER Working Paper No. 17525.

———. 2015. “Efficiency and Policy with Endogenous Learning.” Mimeo.

Bahaj, S., and A. Foulis. 2016. “Macroprudential Policy under Uncertainty.” Working Paper No. 584, Bank of England.

Barlevy, G. 2009. “Policymaking under Uncertainty: Gradualism and Robustness.” *Economic Perspectives* 33 (Q II): 38–55.

Beechey, M. J. 2008. “Lowering the Anchor: How the Bank of England’s Inflation-Targeting Policies Have Shaped Inflation Expectations and Perceptions of Inflation Risk.” Finance and Economics Discussion Series No. 2008-44, Board of Governors of the Federal Reserve System.

Besley, T. 2001. “Political Institutions and Policy Competition.” Technical Report, World Bank.

Bianchi, J., and E. G. Mendoza. 2010. “Overborrowing, Financial Crises and ‘Macroprudential’ Taxes.” NBER Working Paper No. 16091.

Brainard, W. C. 1967. “Uncertainty and the Effectiveness of Policy.” *American Economic Review* 57 (2): 411–25.

Cogley, T., and T. J. Sargent. 2005. “The Conquest of US Inflation: Learning and Robustness to Model Uncertainty.” *Review of Economic Dynamics* 8 (2): 528–63.

Hansen, L. P., and T. J. Sargent. 2001. “Acknowledging Misspecification in Macroeconomic Theory.” *Review of Economic Dynamics* 4 (3): 519–35.

———. 2008. *Robustness.* Princeton University Press.

Jeanne, O., and A. Korinek. 2013. “Macroprudential Regulation versus Mopping Up after the Crash.” NBER Working Paper No. 18675.

King, R. G. 1982. “Monetary Policy and the Information Content of Prices.” *Journal of Political Economy* 90 (2): 247–79.

Korinek, A. 2011. “Systemic Risk-taking: Amplification Effects, Externalities, and Regulatory Responses.” ECB Working Paper No. 1345.

Korinek, A., and A. Simsek. 2014. “Liquidity Trap and Excessive Leverage.” NBER Working Paper No. 19970.

Landes, D. S. 1998. *The Wealth and Poverty of Nations: Why Some Are So Rich and Some So Poor.* London: Little, Brown and Company.

Lorenzoni, G. 2010. “Optimal Monetary Policy with Uncertain Fundamentals and Dispersed Information.” *Review of Economic Studies* 77 (1): 305–38.

Lucas, R. E. 1987. *Models of Business Cycles.* Oxford: Blackwell.

Mukand, S., and D. Rodrik. 2002. “In Search of the Holy Grail: Policy Convergence, Experimentation, and Economic Performance.” NBER Working Paper No. 9134.

Onatski, A., and J. H. Stock. 2002. “Robust Monetary Policy under Model Uncertainty in a Small Model of the U.S. Economy.” *Macroeconomic Dynamics* 6 (01): 85–110.

Orphanides, A., and J. C. Williams. 2007. “Inflation Targeting under Imperfect Knowledge.” *Economic Review* (Federal Reserve Bank of San Francisco) 1: 1–23.

Patton, A. J., and A. Timmermann. 2007. “Testing Forecast Optimality Under Unknown Loss.” *Journal of the American Statistical Association* 102 (480): 1172–84.

Svensson, L. E. 2002. “Monetary Policy and Real Stabilization.” In *Rethinking Policy*, 261–312. Proceedings of an Economic Policy Symposium sponsored by the Federal Reserve Bank of Kansas City, held in Jackson Hole, Wyoming, August 29–31.

Svensson, L. E., and N. M. Williams. 2007. “Bayesian and Adaptive Optimal Policy under Model Uncertainty.” NBER Working Paper No. 13414.

Varian, H. R. 1975. “A Bayesian Approach to Real Estate Assessment.” In *Studies in Bayesian Econometrics and Statistics: In Honor of Leonard J. Savage*, ed. L. J. Savage, S. E. Fienberg, and A. Zellner, 195–208. Amsterdam: North-Holland.

Wald, A. 1950. *Statistical Decision Functions.* New York: John Wiley & Sons.

Weiss, L. 1982. “Information Aggregation and Policy.” *Review of Economic Studies* 49 (1): 31–42.

Wieland, V. 2000. “Monetary Policy, Parameter Uncertainty and Optimal Learning.” *Journal of Monetary Economics* 46 (1): 199–228.

———. 2006. “Monetary Policy and Uncertainty about the Natural Unemployment Rate: Brainard-Style Conservatism versus Experimental Activism.” *B.E. Journal of Macroeconomics* 6 (1, March): 1–34.

Online Appendix to Lower-Bound Beliefs and Long-Term Interest Rates

Christian Grisse,[a] Signe Krogstrup,[b] and Silvio Schumacher[a]
[a]Swiss National Bank
[b]International Monetary Fund

Proof of Proposition 1

The proof closely follows Ruge-Murcia (2006), with a slight modification to allow for a non-zero lower bound. Define $\delta_t \equiv \varepsilon_t/\sigma$ such that $\delta_t \sim N(0,1)$ and

$$r^*_{t+1} = \mathbb{E}(r^*_{t+1}) + \sigma\delta_{t+1},$$

where from (2)

$$\mathbb{E}(r^*_{t+1}) = \alpha + \psi r_t.$$

Then, $r_{t+1} = r^*_{t+1}$ if $r^*_{t+1} > \bar{r}$ or, equivalently, if

$$\delta_{t+1} > \frac{\bar{r} - \mathbb{E}(r^*_{t+1})}{\sigma} \equiv c_{t+1}.$$

Note that c_{t+1} is observed in t. Thus, we have

$$\mathbb{E}(r_{t+1}) = \mathbb{E}(r_{t+1} \mid \delta_{t+1} > c_{t+1})Pr(\delta_{t+1} > c_{t+1}) + \bar{r}Pr(\delta_{t+1} \leq c_{t+1}),$$

where

$$Pr(\delta_{t+1} > c_{t+1}) = 1 - \Phi(c_{t+1})$$

and

$$\begin{aligned}\mathbb{E}(r_{t+1} \mid \delta_{t+1} > c_{t+1}) &= \mathbb{E}(r^*_{t+1}) + \sigma\mathbb{E}(\delta_{t+1} \mid \delta_{t+1} > c_{t+1}) \\ &= (\bar{r} - \sigma c_{t+1}) + \frac{\sigma\phi(c_{t+1})}{1 - \Phi(c_{t+1})},\end{aligned}$$

where the second line uses a result in Maddala (1983, p. 366). Substituting back into the expression for $\mathbb{E}(r_{t+1})$ and simplifying, we get

$$\mathbb{E}(r_{t+1}) = \bar{r} - \sigma c_{t+1}(1 - \Phi(c_{t+1})) + \sigma\phi(c_{t+1}),$$

which is the expression in the text.

Data Sources and Definitions

Table 1. Data Sources and Description

Variable	Source	Datastream Code	Units	Description
Monetary Policy Surprises	Bloomberg and National Central Banks	N/A	bps	Announced policy rate change, less the median from the Bloomberg survey. The survey is conducted among financial analysts, typically on the Friday before the central bank decision.
Interest Rate Swap Yields	Datastream	CCmmY03*	%	Yield on interest rate swap with a variable leg of three months, based on prices as of 5:00 p.m. CET London.
Central Bank Policy Rates				
Reserve Bank of Australia	Datastream	RBACASH	%	Official cash rate
Bank of Canada	Datastream	CN14309	%	Key interest rate (target overnight rate)
Czech National Bank	Datastream	PRREPOR	%	Two-week repo rate
Danmarks Nationalbank	Datastream	DKCTDEP	%	Certificates of deposit
Central Bank of Hungary	Datastream	HNBBASE	%	Base rate
Bank of Japan	Datastream	JPCALLT	%	Uncollateralized call rate
Reserve Bank of New Zealand	Datastream	NZRBCSH	%	Official cash rate
Norges Bank	Datastream	NWFOLIN	%	Sight deposit rate
National Bank of Poland	Datastream	POPRATE	%	Reference rate (seven-day bill rate)
Sveriges Riksbank	Datastream	SDREPOR	%	Repo
Swiss National Bank	Datastream	SWSNBTI	%	Three-month LIBOR target rate
Bank of England	Datastream	LCBBASE	%	Base rate
U.S. Federal Reserve	Datastream	USFDTRG	%	Federal funds target rate
European Central Bank	Datastream	EURODEP	%	Overnight deposit rate

Note: *CC = country code, mm = maturity.

Appendix Tables

Table 2. Panel Regression Results: Robustness Excluding Lower-Policy-Rate Country Observations

	Specification (16)				Specification (17)			
Maturities	**2**	**5**	**7**	**10**	**2**	**5**	**7**	**10**
Δr_t	0.00	−0.06	−0.09	−0.11	0.01	−0.06	−0.07	−0.09
	(0.05)	(0.08)	(0.08)	(0.09)	(0.05)	(0.08)	(0.08)	(0.09)
$r_{i,t-1}$	−1.46***	−1.03*	−0.93*	−0.77				
	(0.44)	(0.57)	(0.55)	(0.56)				
$\Delta r_t \times r_{i,t-1}$	−0.05***	−0.04*	−0.03*	−0.02				
	(0.01)	(0.02)	(0.02)	(0.02)				
$R_{i,t-1}$					−1.40***	−0.86	−0.81	−0.76
					(0.53)	(0.63)	(0.59)	(0.57)
$\Delta r_t \times R_{i,t-1}$					−0.05***	−0.03*	−0.03**	−0.02*
					(0.01)	(0.02)	(0.01)	(0.01)
Constant	−0.90	−2.52*	−3.04*	−3.41*	−0.72	−2.28	−2.59	−2.66
	(0.81)	(1.38)	(1.58)	(1.83)	(0.87)	(1.49)	(1.73)	(2.02)
No. Obs.	137	137	137	137	137	137	137	137
R^2	0.12	0.10	0.11	0.10	0.10	0.08	0.10	0.09

Notes: Observations for countries where the policy rate is already lower than the policy rate reached in the announcing country are excluded. Cluster-robust standard errors are in parentheses. Statistical significance is indicated by *** if $p < 0.01$, ** if $p < 0.05$, and * if $p < 0.1$. $\Delta R_{i,t}$ for the announcing country has been excluded.

Table 3. Panel Regression Results: Robustness Excluding Australia and New Zealand

	Specification (16)				Specification (17)			
Maturities	**2**	**5**	**7**	**10**	**2**	**5**	**7**	**10**
Δr_t	−0.00	−0.05	−0.07	−0.09	0.01	−0.02	−0.02	−0.02
	(0.05)	(0.07)	(0.08)	(0.09)	(0.05)	(0.07)	(0.07)	(0.08)
$r_{i,t-1}$	−1.21***	−1.90***	−2.17***	−2.30***				
	(0.44)	(0.56)	(0.55)	(0.59)				
$\Delta r_t \times r_{i,t-1}$	−0.06***	−0.07***	−0.08***	−0.08***				
	(0.02)	(0.03)	(0.03)	(0.03)				
$R_{i,t-1}$					−1.11**	−1.65***	−1.90***	−2.02***
					(0.49)	(0.55)	(0.49)	(0.51)
$\Delta r_t \times R_{i,t-1}$					−0.06***	−0.07***	−0.08**	−0.07***
					(0.02)	(0.02)	(0.02)	(0.02)
Constant	−1.45**	−2.33*	−2.62*	−2.82*	−1.26*	−1.59	−1.34	−0.87
	(0.73)	(1.20)	(1.41)	(1.66)	(0.76)	(1.14)	(1.35)	(1.62)
No. Obs.	131	131	131	131	131	131	131	131
R^2	0.06	0.10	0.13	0.13	0.04	0.08	0.11	0.10

Notes: Cluster-robust standard errors are in parentheses. Statistical significance is indicated by *** if $p < 0.01$, ** if $p < 0.05$, and * if $p < 0.1$. $\Delta R_{i,t}$ for the announcing country has been excluded.

Table 4. Panel Regression Results: Robustness Excluding Australia, New Zealand, Poland, and Hungary

	Specification (16)				Specification (17)			
Maturities	**2**	**5**	**7**	**10**	**2**	**5**	**7**	**10**
Δr_t	−0.00	−0.04	−0.06	−0.08	0.03*	0.01	0.01	0.01
	(0.04)	(0.06)	(0.07)	(0.08)	(0.02)	(0.03)	(0.02)	(0.03)
$r_{i,t-1}$	−1.37	−2.56***	−2.64***	−2.62**				
	(0.99)	(0.97)	(0.94)	(1.01)				
$\Delta r_t \times r_{i,t-1}$	−0.10*	−0.14**	−0.13**	−0.12*				
	(0.06)	(0.06)	(0.06)	(0.06)				
$R_{i,t-1}$					−1.12	−1.84	−1.92	−2.02*
					(1.22)	(1.25)	(1.23)	(1.21)
$\Delta r_t \times R_{i,t-1}$					−0.10	−0.10	−0.10	−0.09
					(0.08)	(0.08)	(0.08)	(0.08)
Constant	−1.38**	−2.23*	−2.57*	−2.79*	−1.11*	−1.33	−1.20	−0.75
	(0.67)	(1.16)	(1.38)	(1.63)	(0.57)	(0.85)	(0.97)	(1.10)
No. Obs.	107	107	107	107	107	107	107	107
R^2	0.05	0.08	0.08	0.08	0.04	0.07	0.07	0.07

Notes: Cluster-robust standard errors are in parentheses. Statistical significance is indicated by *** if $p < 0.01$, ** if $p < 0.05$, and * if $p < 0.1$. $\Delta R_{i,t}$ for the announcing country has been excluded.

Table 5. Panel Regression Results: Robustness to Omitted-Variable Lower Bound

	Specification (16)				Specification (17)			
Maturities	**2**	**5**	**7**	**10**	**2**	**5**	**7**	**10**
Δr_t	0.02	−0.04	−0.07	−0.09	0.02	−0.03	−0.05	−0.06
	(0.04)	(0.07)	(0.08)	(0.09)	(0.04)	(0.07)	(0.08)	(0.08)
$\tilde{r}_{i,t-1}$	−1.18***	−1.07**	−1.04**	−0.94**				
	(0.38)	(0.45)	(0.42)	(0.42)				
$\Delta r_t \times \tilde{r}_{i,t-1}$	−0.05***	−0.04**	−0.04***	−0.03**				
	(0.01)	(0.01)	(0.01)	(0.01)				
$\tilde{R}_{i,t-1}$					−1.10**	−0.93*	−0.95**	−0.96**
					(0.44)	(0.52)	(0.48)	(0.46)
$\Delta r_t \times \tilde{R}_{i,t-1}$					−0.05***	−0.03**	−0.03***	−0.03***
					(0.01)	(0.01)	(0.01)	(0.01)
Constant	−0.77	−1.84	−2.25	−2.56	−0.68	−1.58	−1.68	−1.55
	(0.82)	(1.28)	(1.47)	(1.71)	(0.91)	(1.42)	(1.62)	(1.85)
No. Obs.	155	155	155	155	155	155	155	155
R^2	0.10	0.11	0.12	0.12	0.09	0.10	0.11	0.10

Notes: Cluster-robust standard errors are in parentheses. Statistical significance is indicated by *** if $p < 0.01$, ** if $p < 0.05$, and * if $p < 0.1$. $\Delta R_{i,t}$ for the announcing country has been excluded. $\tilde{r}_{i,t-1}$ denotes the difference between $r_{i,t-1}$ and the minimum of policy rates observed across the countries in the sample and up to t. A similar definition holds for $\tilde{R}_{i,t-1}$.

Appendix Figures

Figure 1. Response of Two-Year Yields on Interest Rate Swaps around the Riksbank's Policy Announcement of February 12, 2015

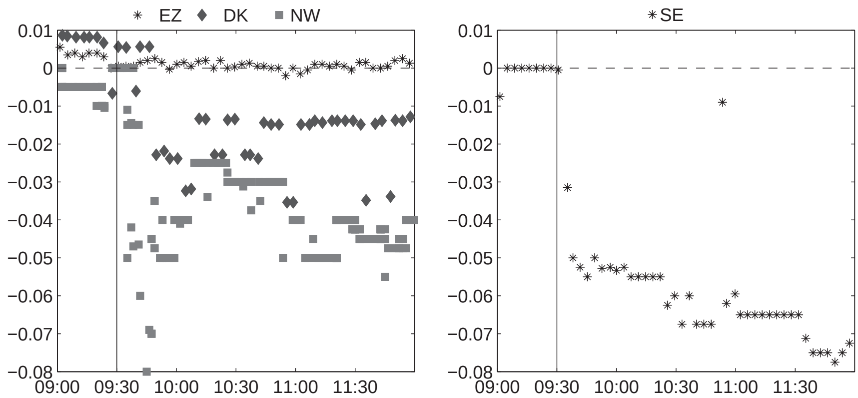

Notes: The cut was unexpected and included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

Figure 2. Response of Two-Year Yields on Interest Rate Swaps around the Riksbank's Policy Announcement of March 18, 2015

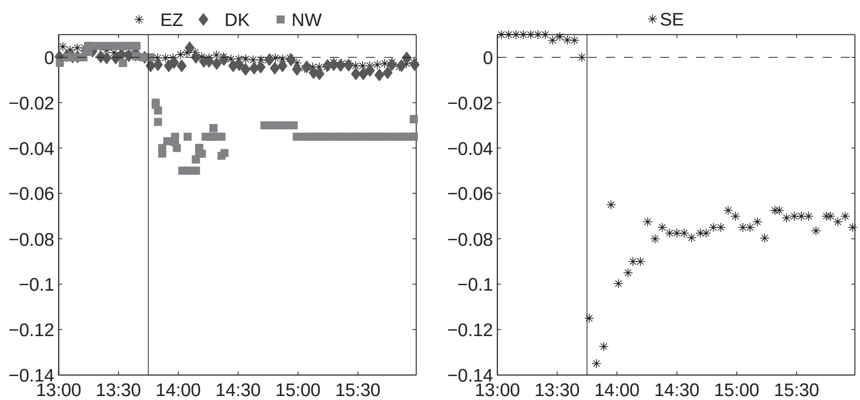

Notes: The cut was unexpected and included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

Figure 3. Response of Two-Year Yields on Interest Rate Swaps around the Riksbank's Policy Announcement of July 2, 2015

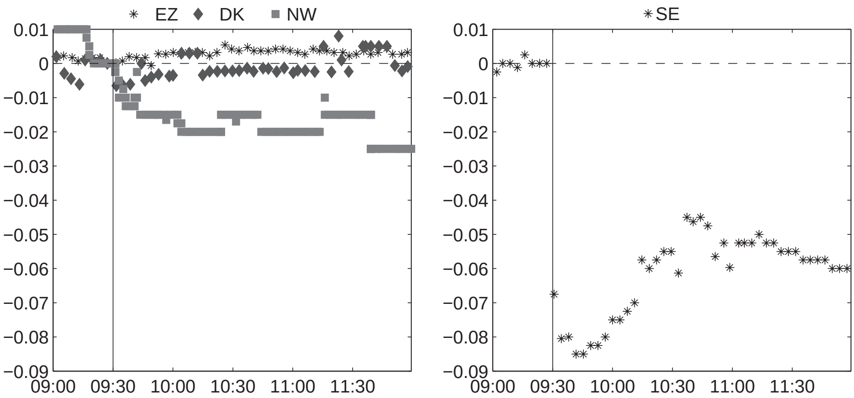

Notes: The cut was unexpected and included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

Figure 4. Response of Two-Year Yields on Interest Rate Swaps around the Riksbank's Policy Announcement of February 11, 2016

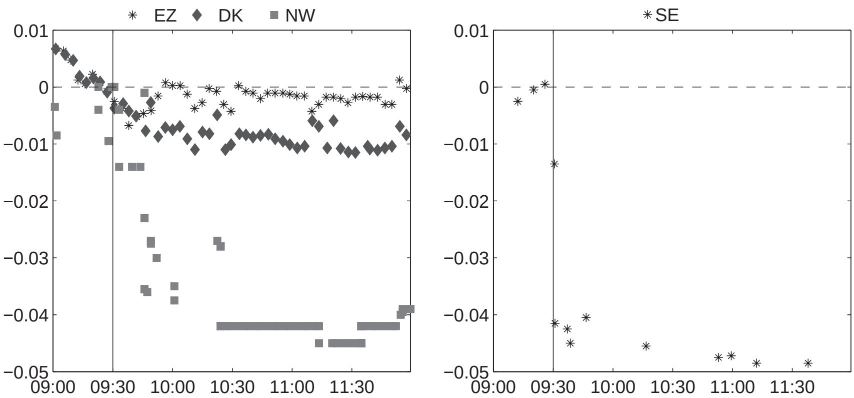

Notes: The cut was partly unexpected and the unexpected component is included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

Figure 5. Response of Two-Year Yields on Interest Rate Swaps around the Danmarks Nationalbank's Policy Announcement of January 19, 2015

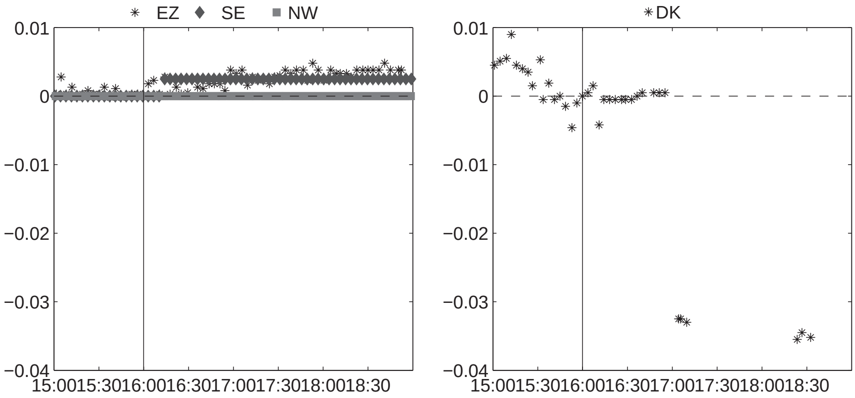

Notes: The cut was unexpected and included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

Figure 6. Response of Two-Year Yields on Interest Rate Swaps around the Danmarks Nationalbank's Policy Announcement of January 22, 2015

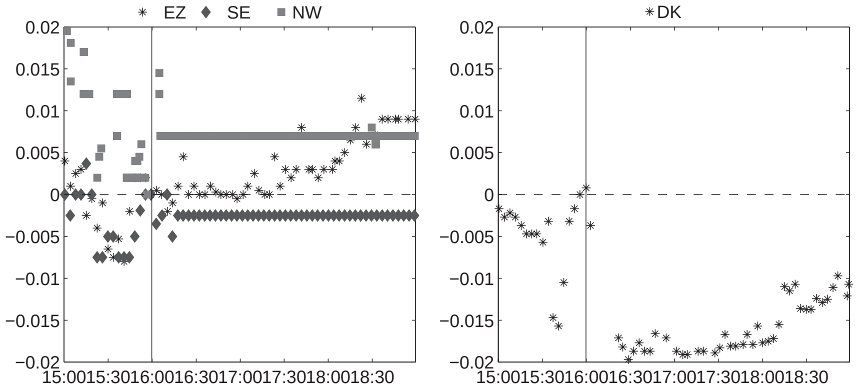

Notes: The cut was unexpected and included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

Figure 7. Response of Two-Year Yields on Interest Rate Swaps around the Danmarks Nationalbank's Policy Announcement of January 29, 2015

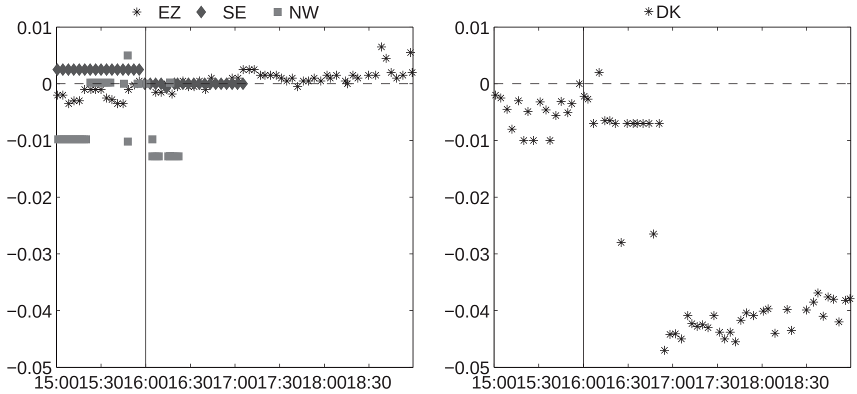

Notes: The cut was unexpected and included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

Figure 8. Response of Two-Year Yields on Interest Rate Swaps around the Danmarks Nationalbank's Policy Announcement of February 5, 2015

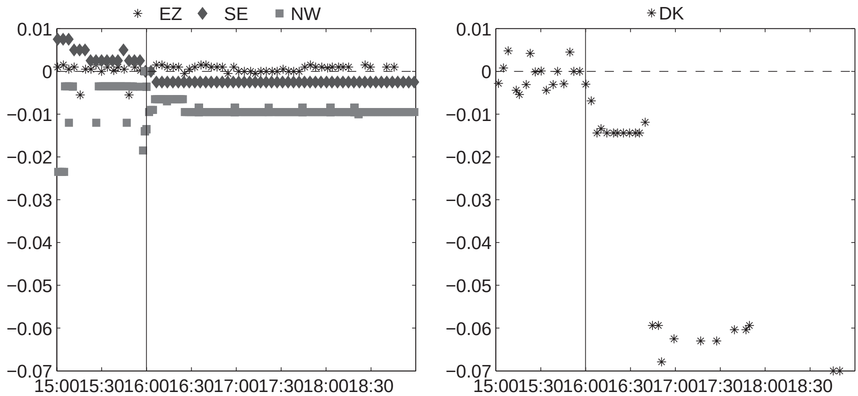

Notes: The cut was unexpected and included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

References

Maddala, G. S. 1983. *Limited-Dependent and Qualitative Variables in Econometrics*. Econometric Society Monograph No. 3. Cambridge: Cambridge University Press.

Discussion of “Macroprudential Policy under Uncertainty”*

François Gourio
Federal Reserve Bank of Chicago

1. Introduction

Following the global financial crisis of 2007–08, policymakers around the world have made it a priority to avoid another occurrence of financial crisis. New, tighter regulations of banks and other financial intermediaries, but also mortgage and other credit markets, hold the promise of reducing both the likelihood and severity of a financial crisis. More controversially, some countries have adopted time-varying macroprudential tools, such that these regulations can be changed according to economic conditions.[1]

Of course, setting the appropriate level of macroprudential “tightness” is challenging, even more so if it has to be varied over time. Policymakers face all the usual difficulties in measuring in real time the relevant economic and financial variables, weighting complicated trade-offs to arrive at an optimal policy, then communicating it and implementing it. On top of that, the effect of the tools is extremely uncertain because these tools are new and untested. To be sure, there is accumulating evidence on the relation between credit, growth, and financial risk (e.g., Schularick and Taylor 2012). However, this evidence is still preliminary and the estimated

*The views expressed here are mine only and do not necessarily reflect those of the Federal Reserve Bank of Chicago or the Federal Reserve System. Author contact: Federal Reserve Bank of Chicago, 230 S. Lasalle St, Chicago IL 60604. E-mail: francois.gourio@chi.frb.org.

[1]Perhaps even more controversially, some policymakers consider using monetary policy tools such as the level of short-term nominal interest rates to affect credit conditions and hence ultimately the risk of financial crisis. See Stein (2013), Ajello et al. (2016), Svennson (2016), and Gourio, Kashyap, and Sim (2017) for discussion and analysis.

relationships may be unstable.[2] The shift in policy regime also means that existing empirical evidence may be less relevant because of the Lucas critique.

In this context, the current paper is a welcome and timely addition. The key contribution is to explain very clearly under which circumstances policymakers should take into account these various forms of uncertainty when formulating policy. The paper builds on a well-established normative literature that studies how uncertainty should affect monetary policy. The starting point for this literature is the certainty equivalence principle, which states that if an agent (the policymaker) is choosing an action to minimize a quadratic expected loss subject to linear constraints and additive shocks, then higher uncertainty about the effect of shocks has no effect on the optimal action: the policymaker takes the same action as if the shock was known to be zero. Of course, higher uncertainty is bad: it increases the expected loss, but it does not affect the optimal action.

Following up on the previous literature, the paper shows that any deviation from these assumptions can generate a deviation from certainty equivalence, i.e., optimal policy is affected by the amount of uncertainty. Clearly, if the loss function is not quadratic—for instance, it is asymmetric—or if the economy is not described by a set of linear equations, the certainty equivalence principle will be violated, though the direction of the bias will depend of course on the exact asymmetry. Also, following Brainard (1967), uncertainty about parameters can lead to attenuation bias; however, as has been recognized for some time, this result is not general and it is easy to build counterexamples where parameter uncertainty lead to a more aggressive policy (see, for instance, Söderström 2002).

Many results discussed in the paper are drawn from literature, but the authors add a very clear exposition and a thoughtful discussion of the practical relevance for macroprudential policy. The authors also propose a new "signaling channel" through which the regulator may use his macroprudential tool not only because it

[2]There is a long history of breakdown of empirical relationships between macroeconomic variables and financial variables, and they may not be informative in real time (e.g., Edge and Meisenzahl 2011), due in part to financial innovations that make consistent financial measurement challenging.

affects the actual risk but also because agents infer from its use that the central bank sees a high level of financial fragility. This allows policy to be more powerful and hence be used more sparingly. An interesting question is whether this allows the policy to be also less distortive (i.e., are signals "cheaper" than actions?).

One potential limitation of the paper is that it studies the policy problem in a simple static linear-quadratic framework. The advantage is the clarity, but it comes at the cost of abstracting from several issues. First, it is not quite clear if some results apply to the *level* of policy variable or to the *change*. For instance, the policy question may be whether to lower the capital requirement from 12 percent to 10 percent. In this context, what does "inaction" or "attenuation bias" mean? Getting a smaller ratio may be less active, but the change in itself is an action.

Second, the static framework abstracts from issues of expectations that may be important, such as these found in the study of monetary policy. For instance, suppose that a regulator adopts a more reactive macroprudential policy rule. This would create expectations by banks of changes in capital requirements in the future. Since banks do not like to have to raise capital quickly, this rule would likely change the behavior of banks which might keep larger buffers (above the capital requirement). These kind of implications are absent from the analysis.

My third issue with the framework has to do with the linear-quadratic assumptions. This is a natural first step, and these assumptions have been commonly used to study monetary policy. However, the framework may be substantially more appropriate for the stabilization problem inherent in monetary policy—i.e., trading off the variance of output against the variance of inflation. It might not adequately capture the trade-offs relevant for prudential policy. For instance, very high financial stability is assumed to be "bad," which on its face seems difficult to understand. I believe it is more natural to think of macroprudential policy as involving a trade-off between the likelihood of crisis with a steady-state cost to the level of output in the economy. Hence, instead of a variance-variance trade-off, one has to consider a mean-variance trade-off. In the rest of this discussion I provide a few simple examples of how one might start thinking about this trade-off.

2. Macroprudential Policy as a Mean-Variance Trade-Off

I first discuss the case with known parameters, then consider the case with unknown parameters.

2.1 The Case with Known Parameters

For simplicity, start with a static framework. Agents have constant relative risk aversion (CRRA) utility over aggregate consumption,

$$U(C) = \frac{C^{1-\gamma}}{1-\gamma}$$

and aggregate consumption is uncertain. To start, assume that it is log-normally distributed:

$$\log C \hookrightarrow N\left(\mu - \frac{\sigma^2}{2}, \sigma^2\right)$$

where μ is the log of the mean of consumption, i.e., $E(C) = e^{\mu}$.[3] Suppose the regulator (government) can take some action that will affect both the mean and the variance of consumption, such as tighter bank capital requirements. Such a measure may reduce variance at the risk of lowering the mean. One can capture this in a reduced-form way by assuming a relationship

$$\mu = \mu_0 + \mu_1 \sigma,$$

where $\mu_1 < 0$ captures the slope of the trade-off.

If the parameters (μ_0, μ_1) describing the trade-off are known, the optimal policy involves choosing a risk that maximizes expected utility, i.e.,

$$\max_{\sigma} E\left(U(C)\right).$$

[3] I incorporate the standard Jensen adjustment, i.e., the mean of log consumption is $\mu - \sigma^2/2$, which ensures that the level of risk σ does not affect the mean of consumption.

It is easy to calculate this expected utility, which turns out to be equal to

$$\max_{\sigma} U\left(\mu_0 - \mu_1\sigma - \gamma\frac{\sigma^2}{2}\right),$$

so that the term inside is the certainty equivalent of consumption. Maximizing expected utility involves then maximizing the certainty equivalent:

$$\max_{\sigma}\left(\mu_0 - \mu_1\sigma - \gamma\frac{\sigma^2}{2}\right).$$

The solution is characterized by the first-order condition which leads to

$$\sigma^* = \frac{-\mu_1}{\gamma},$$

so that higher risk aversion γ will lead to a choice of a lower variance σ^*, while a steeper trade-off (a higher absolute value of μ_1) will lead to a higher risk since reducing risk is more costly.

2.2 The Case with Unknown Parameters

What happens if, on the other hand, the parameters (μ_0, μ_1) are not known with certainty? I will solve this in the simple case where the decisionmaker has a prior normal distribution over each of these parameters. It turns out that the Brainard logic presented by the authors applies in this context,[4] but as we will see, the conclusion might have a slightly different interpretation.

First, to make sure that an increase in uncertainty about μ_0 or μ_1 has no direct effect on expected consumption, I will assume that

$$\mu = \widetilde{\mu_0} + \widetilde{\mu_1}\sigma + k(\sigma),$$

[4]Essentially, the combination of CRRA utility, log-normally distributed endowment, and normal prior yields a quadratic problem similar to the one studied by the authors.

where $k(\sigma)$ is chosen so that $E(C) = e^{E(\mu_0)-E(\mu_1)\sigma}$ does not depend on $Var(\mu_0)$ or $Var(\mu_1)$.[5] Log consumption is still assumed to be distributed with mean $\mu - \frac{1}{2}\sigma^2$ and variance σ^2.

Calculating expected utility (where the expectation is taken now both over the shock and the prior distribution over the parameters) yields

$$E\left(U(C)\right) = U\left(E(\mu_0) - \frac{\gamma}{2}Var(\mu_0) - E(\mu_1)\sigma - \gamma\frac{\sigma^2}{2}Var(\mu_1) - \gamma\frac{\sigma^2}{2}\right)$$

so that the optimal policy problem is to maximize the certainty equivalent:

$$\max_{\sigma} E(\mu_0) - \frac{\gamma}{2}Var(\mu_0) - E(\mu_1)\sigma - \gamma\frac{\sigma^2}{2}Var(\mu_1) - \gamma\frac{\sigma^2}{2}$$

with first-order condition leading to the optimal choice:

$$\sigma^* = \frac{-E(\mu_1)}{\gamma(1 + Var(\mu_1))}.$$

We see that the same two conclusions as in the Brainard model hold here: first, uncertainty about μ_0 is irrelevant for the optimal decision—the term $Var(\mu_0)$ does not appear in the optimal policy choice σ^*, but it does reduce the expected utility. Second, uncertainty about μ_1 leads to a choice of a lower variance σ^* than the policymaker would choose if he had full information about the parameters. One might think that is a standard Brainard result, and indeed it is, except for the interpretation—a low σ^* means, in practice, a very tight macroprudential policy, so the policymaker is actually using his tools aggressively despite the high uncertainty with which they affect the average consumption of the economy. But as discussed above, this reflects that the Brainard result does not hold under all conditions—it depends if a more aggressive policy increases or reduces risk.

[5]This requires that

$$k = -\frac{1}{2}Var(\mu_0) - \frac{1}{2}\sigma^2 Var(\mu_1),$$

simply reflecting the standard Jensen adjustment.

The result above relies on many assumptions, including the functional form linking μ and σ, as well as the utility function and the log-normal assumption. It should be clear that altering these assumptions could easily lead to a reversal of the result.

2.3 Modeling Financial Crises as Rare Disasters

The modeling so far has assumed that macroprudential policy aims to stabilize the economy by reducing the variance of consumption. Macroprudential policy is, however, mostly focused on reducing the likelihood or severity of rare financial crises. One might worry then that the modeling adopted above, while tractable, does not capture the key trade-offs appropriately.

To illustrate this, I consider the kind of processes studied in Barro (2006) or Gourio (2012): consumption may have either a high level $\overline{C}$ or a low level $\overline{C}(1-b)$. The first case arises with probability $1-p$ and the second one with probability p. The scalar p can be thought of as the probability of a financial crisis, and b measures the share of consumption lost in a financial crisis. In this case, it is easy to see that expected utility is

$$EU(C) = \frac{\overline{C}^{1-\gamma}}{1-\gamma}\left(1-p+p(1-b)^{1-\gamma}\right).$$

If one assumes that there is a trade-off between the level of consumption and the probability of crisis, for instance,

$$\log \overline{C} = d_0 + d_1 p,$$

where the parameters (d_0, d_1) reflect the trade-off between the level of output and the probability of financial crisis, then it is straightforward to obtain the optimal p. One could also do the same calculation if there is uncertainty about d_1. Clearly, the functional forms and assumptions will drive the result. The general spirit is quite similar to that of the log-normal example. But the quantitative results might be quite different, as the welfare cost of avoiding rare and large shocks can be quite larger than smoothing out small fluctuations (e.g., Barro 2009; Gourio, Kashyap, and Sim 2017).

2.4 Generalization to Infinite Horizon

For simplicity I have focused on static calculations in these examples, and I have used a simple power utility. The goal of this section is simply to clarify that this comes at no cost of generality, because it is straightforward to generalize these results to an infinitely lived agent with Epstein-Zin preferences, at least as long as consumption growth is iid. To see this, consider a representative agent in which utility (value) is defined recursively as

$$V_t = \left((1-\beta) C_t^{1-1/\psi} + \beta E_t \left(V_{t+1}^{1-\gamma} \right)^{\frac{1-1/\psi}{1-\gamma}} \right)^{\frac{1}{1-1/\psi}},$$

where ψ is the elasticity of intertemporal substitution to consumption and γ measures risk aversion. Assume that this consumer has an initial consumption C_0 and faces macroeconomic risk in the form of a unit-root iid consumption process:

$$\Delta \log C_{t+1} = X_{t+1},$$

where X_{t+1} is an iid random variable. For instance, X_{t+1} could be normally distributed with mean $\mu - \frac{\sigma^2}{2}$ and variance σ, the standard log-normal case. But one could also allow X_{t+1} to take the form of rare financial crises, so that $X_{t+1} = \mu$ with probability $1-p$ and $X_{t+1} = \mu - b$ with probability p.

To find the optimal policy, we must first solve for welfare. It is easy to see in this model that the value-consumption ratio V_t/C_t is constant and given by the solution to

$$v^{1-\frac{1}{\psi}} = 1 - \beta + \beta v^{1-\frac{1}{\psi}} E_t \left(X_{t+1}^{1-\gamma} \right)^{\frac{1-1/\psi}{1-\gamma}}$$

and hence is determined by the moment $E_t(X_{t+1}^{1-\gamma})$, which is the one we studied in the static model. Hence our conclusions should apply in this framework as well.

3. Conclusion

The paper is a clean and lucid analysis of how uncertainty should affect macroprudential policy. It is hard to draw strong conclusions given the various channels and effects highlighted by the authors. Moving beyond the qualitative results will require a credible DSGE model that incorporates financial crises, makes their probability and severity affected by macroprudential policy, and also incorporates the costs of tighter regulation. Such a model would address the key trade-off created by macroprudential policy, namely that lower volatility (lower financial crisis probability) comes at a cost of less capital accumulation and less production.

To conclude, one should also keep in mind that the design of macroprudential policy cannot be isolated from that of the other policies (monetary and fiscal). A natural approach is to let macroprudential policy focus on financial stability and let the other policies offset the indirect effect of macroprudential policy on output and inflation. This might create an additional constraint on monetary and fiscal policies, and transfer the risk from the regulatory to the monetary authority.[6]

References

Ajello, A., T. Laubach, D. López-Salido, and T. Nakata. 2016. "Financial Stability and Optimal Interest-Rate Policy." Working Paper, Board of Governors of the Federal Reserve System.

Barro, R. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics* 121 (3): 823–66.

———. 2009. "Rare Disasters, Asset Prices, and Welfare Costs." *American Economic Review* 99 (1): 243–64.

Brainard, W. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review* 57 (2): 411–25.

Edge, R., and R. Meisenzahl. 2011. "The Unreliability of Credit-to-GDP Ratio Gaps in Real Time: Implications for Countercyclical Capital Buffers." *International Journal of Central Banking* 7 (4, December): 261–98.

[6]For instance, in the spirit of Evans et al. (2015), a tight macroprudential policy might make it more likely that monetary policy is constrained (now or in the future) by the zero lower bound on nominal interest rates, increasing the risk of not reaching its targets and increasing uncertainty for the economy more generally.

Evans, C., J. Fisher, F. Gourio, and S. Krane. 2015. "Risk Management near the Zero Lower Bound." *Brookings Papers on Economic Activity* (Spring): 141–219.

Gourio, F. 2012. "Disaster Risk and Business Cycles." *American Economic Review* 102 (6): 2734–66.

Gourio, F., A. Kashyap, and J. Sim. 2017. "The Tradeoffs in Leaning Against the Wind." Working Paper, Federal Reserve Bank of Chicago.

Schularick, M., and A. Taylor. 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008." *American Economic Review* 102 (2): 1029–61.

Söderström, U. 2002. "Monetary Policy with Uncertain Parameters." *Scandinavian Journal of Economics* 104 (1): 125–45.

Stein, J. 2013. "Overheating in Credit Markets: Origins, Measurement, and Policy Responses." Speech given at the "Restoring Household Financial Stability after the Great Recession: Why Household Balance Sheets Matter" research symposium sponsored by the Federal Reserve Bank of St. Louis, held in St. Louis, Missouri, February 7. Available at https://www.federalreserve.gov/newsevents/speech/stein20130207a.htm.

Svensson, L. E. O. 2016. "Cost-Benefit Analysis of Leaning Against the Wind." NBER Working Paper No. 21902.

Lower-Bound Beliefs and Long-Term Interest Rates*

Christian Grisse,[a] Signe Krogstrup,[b] and Silvio Schumacher[a]
[a]Swiss National Bank
[b]International Monetary Fund

We study the transmission of changes in the believed location of the lower bound to long-term interest rates since the introduction of negative interest rate policies. The expectations hypothesis of the term structure combined with a lower bound on policy rates suggests that the transmission of policy rate changes to long-term interest rates is reduced when policy rates approach this lower bound. We show that if market participants revise downward the believed location of the lower bound, this may reduce long-term yields and increase transmission. A cross-country event study suggests that such effects have been empirically relevant during the recent negative interest rate episode.

JEL Codes: E43, E52.

1. Introduction

Policy interest rates were reduced to near zero in many countries in the aftermath of the global financial crisis. Because zero was broadly believed to be the lower bound on interest rates, further monetary loosening was achieved mainly through unconventional monetary policies, such as quantitative easing. More recently, however, some

*We thank Ran Bi, Hans Dewachter (discussant), James Hamilton (discussant), Emanuel Mönch, Thorvardur Olafsson, Rafael Repullo, Eric Swanson, Andreea Vladu (discussant), and participants at the ECB workshop on negative interest rates, the Bundesbank-OeNB-SNB research workshop, and the Annual International Journal of Central Banking Research Conference 2016 for comments and suggestions. We also thank Rafael Greminger for excellent research assistance. The views in this paper are solely the responsibility of the authors and do not necessarily reflect the views of the Swiss National Bank, the International Monetary Fund, its Management, or its Board. Author contact: christian.grisse@snb.ch, skrogstrup@imf.org, silvio.schumacher@snb.ch.

central banks cut their policy rates into negative territory. Danmarks Nationalbank, the Danish central bank, was a pioneer in cutting its policy rate to −25 basis points in mid-2012.[1] The European Central Bank (ECB) cut its policy rate below zero in mid-2014, which was followed by incremental rate cuts below zero by the Swiss National Bank (SNB), Danmarks Nationalbank, and Sveriges Riksbank. The Bank of Japan cut its main policy rate to −10 basis points in January 2016. The SNB has gone furthest with negative policy rates by introducing negative interest rates of −75 basis points on banks' sight deposit account balances, which have remained in place since January 2015.

As policy rates have been cut below zero, the market's belief about the location of the lower bound on nominal interest rates has also moved into negative territory. The lower bound remains unknown, however. We refer to the level of the policy interest rate that markets broadly believe is the lower bound as the believed lower bound.

Interest rate policy transmits to the real economy through its effect on the broader term structure of interest rates, and in normal times this transmission is not constrained by the lower bound on interest rates. When the policy rate is very low or negative, the distance to the believed lower bound narrows, and room for future downward movements in policy rates shrinks. Ruge-Murcia (2006) uses a very simple two-period term structure model to show that this reduces the transmission of short-term policy rate cuts to long-term interest rates. Ruge-Murcia (2006) and Grisse (2015) present empirical evidence based on U.S. and Japanese data suggesting that the transmission of short-term rates to long-term rates is indeed reduced when policy rates are closer to zero. We argue and provide empirical evidence that if rate cuts below zero simultaneously move the market's believed location of the lower bound, this influences the distribution of possible future short-term policy rates and may thus directly move long-term interest rates and increase policy transmission.

[1] Sweden had a brief period of a negative rate on its deposit facility in 2009, but this rate was not effectively binding for money markets and hence did not transmit to financial markets.

We first theoretically study the transmission of changes in the believed location of the lower bound to long-term interest rates. To motivate our empirical analysis, we extend the theoretical term structure model of Ruge-Murcia (2006) to show that when short-term interest rates are held constant, a drop in the believed lower bound leads to a drop in longer-term interest rates, but only when long-term interest rates are initially constrained by the lower bound on short rates. This mechanism has also been highlighted by Kortela (2016) and Lemke and Vladu (2016) based on much more sophisticated shadow-rate term structure models. The advantage of illustrating this point with the Ruge-Murcia model is that it is analytically very simple and intuitive, though at the cost of simplifying other factors that drive yield-curve variation. Second, we empirically study the predictions of the model by using an event study of a panel of advanced countries. To measure changes in the believed location of the lower bound, we focus on central bank announcements of interest rate cuts into negative territory, which likely departed from conventional wisdom on the location of the lower bound. We conjecture that such announcements changed the perception of the lower bound not only in announcing countries but also across countries more broadly. In other countries, short-term policy rates would have remained unchanged on these occasions. By focusing on the effects of such announcements on other countries' long-term interest rates and relating these effects to the initial level of short-term interest rates across countries, we can identify the effect of changes in the perceived location of the lower bound independently from the effect of policy rate changes and more standard international interest rate spillovers. Our results lend support to the model's prediction.

There are important policy implications of these findings. In the vicinity of the lower bound on interest rates, central bank actions or communication that result in a downward shift in the believed lower bound can reduce longer-term interest rates and thereby achieve monetary policy loosening *independently* of short rate cuts. Such actions or communication can also increase the distance of the policy rate to the believed lower bound, thereby increasing the transmission of future rate cuts to long-term interest rates. Managing market perceptions regarding the location of the lower bound is an important dimension of effective central bank communication when policy rates are low.

Our paper is related to the two previously mentioned papers that investigate the implications of a changing lower bound for monetary policy transmission in the context of shadow-rate term structure models. Based on their model, Lemke and Vladu (2016) use survey expectations on interest rates in the euro area to argue that there were two changes in the market-perceived lower bound prior to 2015: a first change in May 2014 (when some forecasters started to predict an ECB rate cut to −10 basis points at the June 2014 meeting) and a second change in September 2014 (when the ECB's rate cut of 10 basis points took markets by surprise). According to their model, the decline in euro-area yields around the ECB's September 2014 rate cut was exclusively driven by changes in the lower bound. Kortela (2016) also estimates a shadow-rate term structure model using euro-area data, but the lower bound is allowed to vary gradually over the sample. Considering alternative exogenous processes for the lower bound, Kortela (2016) finds that models with a time-varying lower bound fit the data better than constant lower-bound models. The model also suggests that lower-bound changes were the main driver of euro-area two-year yield movements between September 2014 and January 2016. In contrast, the contribution of lower-bound changes to changes in ten-year yields is found to be small. Our paper is the first to consider the effects of changes in the believed location of the lower bound on long-term yields in broader panel data on all five countries that have moved policy rates into negative territory.

The remainder of this paper is structured as follows. The next section presents and extends the Ruge-Murcia (2006) model, which motivates the empirical analysis and identification strategy. The empirical analysis is presented in section 3, and section 4 concludes.

2. Theory

With the purpose of deriving an empirical identification strategy for assessing the relevance of changes in markets' perceptions of the location of the bound, we turn to the very simple two-period term structure model with a lower bound derived in Ruge-Murcia (2006).[2]

[2]Ruge-Murcia (2006) shows that his results also apply to a more general setting with an n-period bond, a more general stochastic process for the shadow

The Ruge-Murcia model is highly intuitive and can be solved analytically, but it nevertheless has sufficient richness to allow us to characterize the lower-bound effects on long-term rates in which we are interested. We extend the Ruge-Murcia model to allow for the lower bound to be different from zero and derive the impact of a change in this lower bound on longer-term interest rates. In the empirical approach, we proxy changes in the believed location of the lower bound by central bank announcements of policy rates moving into negative territory. This gives rise to the empirical identification problem that both changes in short-term interest rates and changes in the believed lower bound affect long-term rates and occur at the same time. In section 3, we use the model to derive the conditions under which the presence of shifts in the believed lower bound can be identified in the data.

The short-term interest rate r_t is constrained by a lower bound,

$$r_t = \max\left(r_t^*, \bar{r}\right), \tag{1}$$

where r_t^* is the "shadow" short rate and $\bar{r}$ is the level of the nominal interest rate that market participants jointly believe is the lower bound. This believed lower bound is non-stochastic, and we treat it as exogenous in the model in order to derive very simple derivatives with respect to this level. We assume that it is equal to or below the period 1 level of the short-term interest rate.

While it is outside the model, we can think of the empirically unobservable believed lower bound as being influenced by observable factors that are believed to influence the location of the actual lower bound. The absolute or "true" lower bound is the level at which a further interest rate cut will cause prohibitive financial upheaval—or the level at which the central bank is simply unwilling to move the policy rate further down, independently of such financial market reactions. Prohibitive financial upheaval, for example, could be a large-scale demand for zero-interest-rate physical cash in exchange for bank deposits or other assets, effectively precluding the transmission of the rate cut to the broader financial markets. A central

short rate, and non-normally distributed shocks. However, in this more general setting, no closed-form solution exists, and the model properties can only be explored numerically.

bank's unwillingness to reduce the policy rate in the absence of such prohibitive financial upheaval can be driven by cultural, communication, or political economy considerations. Negative interest rates have turned out to be unpopular and difficult to understand for the broader population, giving rise to difficult challenges with communication. Central banks' assessment of the degree to which banks are willing to transmit negative interest rates may also matter. Banks' reactions to rate cuts below zero matter for the expected effectiveness of the transmission of a policy rate cut through the bank lending channel, which may change the cost-benefit analysis of a rate cut for central banks.[3] They also matter for the expected transmission of negative rates to retail deposits, which may increase the risk of a large-scale shift into cash directly; see also Danthine (2016). The design features of negative policy rates that matter for how bank profits are affected, such as the central bank operating framework and possible exemptions to the negative interest rate on central bank reserves, may thus also matter for the location of the "true" lower bound.

The discussion suggests that the location of the lower bound is multifaceted, country specific, perhaps time varying, and very difficult to estimate. It is as of yet unobserved, but the market forms a belief about the location of this lower bound, depending on perceptions of financial market structure and preparedness, the central bank operating framework, and cultural and political considerations. As long as the believed lower bound is below the prevailing policy interest rate level and the true lower bound is unobserved, it is the believed lower bound that matters for expectations about future interest rates.

The shadow short-term interest rate depends on past short-term rates,

$$r_t^* = \alpha + \psi r_{t-1} + \varepsilon_t, \tag{2}$$

where $\psi > 0$, such that interest rates are positively autocorrelated, and $\varepsilon_t \sim N(0, \sigma^2)$. The model is reduced to the setting of

[3]One way to think about this bank transmission channel is the reversal rate, as in Brunnermeier and Koby (2016).

Ruge-Murcia (2006) if $\bar{r} = 0$. When $r_t^* < \bar{r}$, the shadow rate is unobserved, and market participants have to form expectations about r_t^* conditional on variables observed up to and including period t.[4]

Finally, the yield of a two-period bond R_t is determined by the expectations hypothesis,

$$R_t = \frac{1}{2}\left(r_t + \mathbb{E}(r_{t+1})\right) + \theta_t, \tag{3}$$

where θ_t is a liquidity and term premium, which is assumed to be uncorrelated with ε_t.

To solve the model and express R_t as a function of r_t, we have to compute $\mathbb{E}(r_{t+1})$. This is done in the following proposition, which is proved in Ruge-Murcia (2006) for the case where $\bar{r} = 0$.

PROPOSITION. *The expected future short rate is given by*

$$\mathbb{E}(r_{t+1}) = \bar{r} + (\alpha + \psi r_t - \bar{r})\left(1 - \Phi(c_{t+1})\right) + \sigma\phi(c_{t+1}), \tag{4}$$

where

$$c_{t+1} \equiv \frac{\bar{r} - \mathbb{E}(r_{t+1}^*)}{\sigma} = \frac{\bar{r} - (\alpha + \psi r_t)}{\sigma} \tag{5}$$

and $\Phi(\cdot)$ and $\phi(\cdot)$ denote the standard normal cumulative and density functions.

Proof. See the online appendix, at http://www.ijcb.org.

Here, c_{t+1} can be understood as a summary measure of the degree to which the policy rate will be constrained by the believed lower bound in period 2. Substituting (4) into (3) gives the solution for the long-term bond yield,

$$R_t = \frac{r_t + \bar{r}}{2} + \frac{1}{2}(\alpha + \psi r_t - \bar{r})\left(1 - \Phi(c_{t+1})\right) + \frac{\sigma}{2}\phi(c_{t+1}) + \theta_t. \tag{6}$$

[4]One would expect that $\mathbb{E}(r_{t+1}) \geq \bar{r}$. From (4), a sufficient condition for this to be the case is that $c_{t+1} \leq 0$, or equivalently, $\mathbb{E}(r_{t+1}^* \mid I_t) = \alpha + \psi r_t \geq \bar{r}$. Since $r_t \geq \bar{r}$, we require $\alpha/(1-\psi) \geq \bar{r}$. Assuming that $|\psi| < 1$, this inequality suggests that the unconditional mean of r_t^* needs to be above the lower bound on r_t. However, this condition is not relevant any more if, as in the general model in Ruge-Murcia (2006), additional exogenous variables are added to the shadow-rate process (2).

The main results of Ruge-Murcia (2006) immediately follow from the first and second derivatives of (6). Not surprisingly, we have

$$\frac{\partial R_t}{\partial r_t} = \frac{1}{2} + \frac{\psi}{2}\left(1 - \Phi(c_{t+1})\right) > 0. \tag{7}$$

However, the response of long yields to short rate movements is stronger if interest rates are higher:

$$\frac{\partial^2 R_t}{\partial r_t^2} = \frac{\psi^2}{2\sigma}\phi(c_{t+1}) > 0. \tag{8}$$

This result implies that as nominal interest rates approach the lower bound, the response of longer-term yields to short rate changes declines. This non-linearity arises because of the effect of r_t on $\mathbb{E}(r_{t+1})$, through the expectations hypothesis of the term structure. The intuition for these results is similar to that for the "honeymoon" effect (the non-linear relationship between the exchange rate and its fundamentals) in Krugman's (1991) target zone model. Suppose that market participants think that future positive and negative interest rate shocks are equally likely. In addition, suppose that short-term rates are at 0.5 percent, with a lower bound of zero. Then, a future shock to the shadow short rate of +1 percentage points would raise short rates to +1.5 percent, while a shock of –1 percentage points would only lower short rates to zero. Because market participants anticipate that future short rate shocks are constrained by the lower bound on nominal interest rates in this way, expected future short rates and hence the yield curve are affected even when short rates are still above the lower bound. When short rates are closer to the lower bound, a short rate decline will produce a smaller downward shift in expected future short rates and therefore also a smaller effect on long-term yields.

We have thus far assumed that $\bar{r}$ is fixed. Now consider a surprise change in the lower bound—for example, due to a central bank announcement: in t, the lower bound unexpectedly changes, but market participants assume that $\bar{r}$ will remain fixed at the new value from t onward. In this exercise, we keep the current short rate r_t unchanged, as would be the case if the central bank communicates that policy rates are possible at levels below the current believed location of the lower bound—or if the revision in market views about the lower bound reflects the actions of a foreign central bank. From (6), we have

$$\frac{\partial R_t}{\partial \bar{r}} = \frac{1}{2}\Phi(c_{t+1}) > 0. \tag{9}$$

It is intuitive that this expression is positive, for the same reasons as outlined in the previous paragraph.[5] Furthermore, as expected, $\partial R_t/\partial \bar{r}$ is increasing in c_{t+1}, i.e., a change in market views of the lower bound matters most in a situation where short rates are already constrained. From the definition of c_{t+1}, it follows that

$$\frac{\partial^2 R_t}{\partial \bar{r} \partial r_t} = -\frac{\psi}{2\sigma}\phi(c_{t+1}) < 0. \tag{10}$$

Figure 1 depicts the relationship between the short-term policy rate and the long-term interest rate based on a simulation of the model, for two different levels of the lower bound. The black line reflects the relationship given a believed lower bound of zero, whereas the gray dashed line reflects a believed lower bound below zero. The black line has more curvature around zero, reflecting a higher degree of constraint emanating from the assumed belief that interest rates cannot go below zero in the second period. When the believed lower bound shifts downward into negative territory, the long-term interest rate shifts down given the policy rate, and this downward shift is greater when the policy rate is closer to the initial location of the believed lower bound.

In our empirical identification strategy, we utilize the finding that the closer the policy rate to the perceived lower bound, the weaker the effect of changes in short-term interest rates on long-term rates, but the stronger the effect of changes in the believed lower bound on long-term interest rates.

3. An Event Study of Changes in Lower-Bound Beliefs

The model shows that when the believed location of the lower bound decreases, this reduces long-term interest rates to the extent that

[5]The same result is derived in the more sophisticated term structure model of Lemke and Vladu (2016).

Figure 1. Illustration of the Link between the Short-Term Interest Rate and the Long-Term Interest Rate in the Two-Period Ruge-Murcia (2006) Model

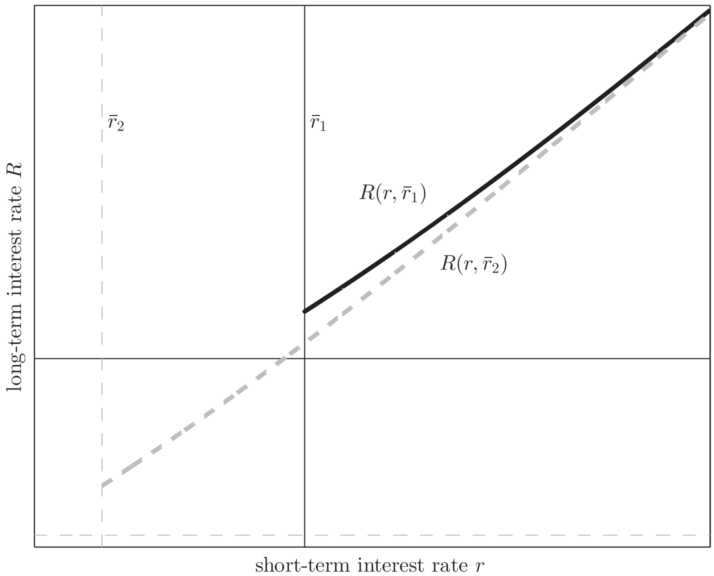

Notes: Parameters: $\alpha = 0, \psi = 0.95, \sigma = 1$. The solid line assumes that the lower bound on the policy rate is $\bar{r}_1 = 0$ percent, while the dashed line assumes a lower bound of $\bar{r}_2 = -0.75$ percent.

these are already constrained by the lower bound. According to (9), the model predicts that the decrease in the long-term interest rate associated with a decrease in the believed location of the lower bound is positively related to the proximity of the lower bound. We empirically assess this prediction by using an event-study approach for a panel of fourteen advanced economies. Empirically, we proxy changes in the believed lower bound by unexpected central bank announcements of policy rate cuts below zero. To study the effects in a panel of countries, we conjecture that unexpected central bank cuts below zero have affected the believed lower bound globally and not just in the interest-rate-cutting country. We then study how long-term yields have reacted across countries to such announcements. There are a number of identification challenges with this approach, as well as event-study design issues, and we first discuss how we address these.

3.1 Identification

To identify changes in the believed location of the lower bound, we assume that central bank announcements of interest rate cuts that departed from conventional wisdom regarding the location of the lower bound are likely to have moved the perceived location of the lower bound down. We also assume that such announcements changed the believed location of the lower bound in countries that did not simultaneously change their short-term policy rates. By focusing on the reaction of yields in non-announcing countries, the theoretical framework allows us to identify the presence of changes in the perceived location of the lower bound on long-term interest rates. To do this, we first make the following assumption:

ASSUMPTION 1. *The cross-country variation in the effect on the country-specific believed lower bound in non-announcing countries is independent of the interest rate level in those countries.*

Practically, we think of the change in the believed lower bound in country i at time t, $\Delta\bar{r}_{i,t}$, as being proportional to the change in the policy rate in the announcing country, Δr_t:

$$\Delta\bar{r}_{i,t} = \gamma\Delta r_t + \upsilon_{i,t} \qquad (11)$$
$$\gamma > 0,$$

where $\upsilon_{i,t}$ are independently distributed errors that capture country-specific factors influencing the strength of the revision in the believed lower bound. Expression (11) implies that the change (though not necessarily the level) of the believed lower bound across non-announcing countries i in response to a negative interest rate announcement in period t is proportional to the unexpected part of the cut in the announcing country. We do allow for this effect to be smaller in non-announcing countries than the effect on the believed lower bound in the announcing country, however (e.g., if $\gamma < 1$). A rate cut below zero could affect the believed lower bound more in the rate-cutting country than in other countries if the believed lower bound also captures an assessment of how far a particular central bank is ready to cut rates, for example, based on the degree of conservatism of the institution and country or the country-specific state and readiness of the financial sector with respect to negative rates.

Assumption 1 means that the effect on the believed lower bound in other countries is independent of the prevailing interest rate level in those other countries (i.e., the country-specific error term is independent of the interest rate level). We think that this assumption is reasonable. Possible cross-country variation in the effect of a rate cut in other countries' believed lower bound is likely to depend notably on country-specific features as mentioned above, but such characteristics in turn are unlikely to be structurally linked with the interest rate level. Assumption 1 is actually stronger than required for our purposes. In fact, we require only that $\Delta\bar{r}_{i,t}$ is not positively related to $r_{i,t}$ (i.e., that beliefs are not revised downward more in countries with higher interest rates). One could perhaps argue that for behavioral reasons, updates to the believed lower bound occur only (or are larger) if the interest rate is sufficiently low. For example, in a rational inattention setting where belief updating is costly, investors may be more likely to contemplate the implications of foreign announcements for the domestic lower bound if interest rates are sufficiently close to the lower bound already. In this case, the strength of belief revisions would depend negatively on $r_{i,t}$. This does not affect our identification strategy, which relies on the idea that changes in lower-bound beliefs are the only mechanism that generates a negative correlation between the strength of changes in long-term yields and the policy rate changes.

The selection of the relevant set of belief-changing central bank announcements is not straightforward. There are no data or surveys that provide systematic information on changes in the market's belief about the location of the lower bound. It is very difficult to come up with objective criteria for selecting announcements that are likely to be associated with changes in the believed lower bound. Identification hence has to rely on a certain amount of judgment. To remain as general and objective as possible but with the risk of selecting announcements that are not associated with changes in the believed location of the lower bound, our baseline includes all central bank announcements of leading policy rate cuts into, or further into, negative territory that were unexpected.[6] This event-selection strategy may increase the amount of noise in the independent variable, which could give rise to attenuation bias. This would stack the odds against finding significant effects, suggesting that the effects may be even stronger than those suggested by our estimates. Table 1 presents all the announcements made to date of leading policy rate cuts in negative territory. Column 6 contains the part of the announced cut that was unexpected, as measured by Bloomberg's survey of financial market participants. We thus select the twelve announcements that contained an element of surprise according to the Bloomberg survey.

As it is unlikely that all these announcements had similar effects on the believed lower bound, we also consider alternative announcement selections. For example, later cuts into negative territory could have been less belief changing than early ones, where interest rates went below levels never seen before. As a robustness check of our baseline, we therefore also consider the subset of "new global lows" (i.e., interest rate cuts that moved a policy rate below levels previously seen). This subset comprises the two announcements by the SNB, as earlier announcements by Danmarks Nationalbank that first brought interest rates into negative territory were anticipated. We also consider a sample that excludes country observations where the policy rate of the country in question was already lower than that reached by the announcing country. The drawback of these two subsets, particularly the former, is that the number of observations is reduced.

A second identification challenge is that policy rate cuts may also affect long-term interest rates through the normal channels of the term structure of interest rates. As discussed above and as shown by the model, a policy rate cut directly affects long-term yields in the

[6] We select only cuts below zero for the policy rate that transmit to money markets and do not consider cases where interest rates on central bank deposit facilities were cut into negative territory when the leading policy rate has remained positive. The first announcement by Riksbank in 2009 of a negative interest rate on its deposit facility did not effectively bind or transmit to money markets and is hence unlikely to have changed perceptions of how much interest rates could be cut below zero. More recently, both Norway and Hungary cut the rate on their deposit facilities below zero while keeping the leading policy rate positive. We do not consider these announcements to be relevant either.

Table 1. Negative Interest Rate Announcements

Country	Date, Time	Policy Rate Type	Prev. Pol. Rate (%)	Pol. Rate Change (bps)	Surpr. Comp.[a] (bps)	Gov. Yield Change[b]	Additional Policy Announcements
Switzerland	Jan. 15, 2015 10:31	Sight Deposits	−0.25	−50	−50	−8	Discontinue minimum exchange rate. Lower target range for the three-month LIBOR from (–0.75% to –0.25%) to (–1.25% to –0.25%).
Switzerland	Dec. 18, 2014 08:00	Sight Deposits	0.00	−25	−25	−8	Lower target range for the three-month LIBOR to (–0.75% to –0.25%). Exemption threshold for negative rates.
Euro Area	Mar. 10, 2016 13:45	Deposit Facility	−0.3	−10	0	8	Main refinancing rate decreased by 5 bps to 0%, lending facility decreased by 5 bps to 0.25%. Monthly asset purchase program expanded to 80 billion. Launch new four-year targeted longer-term refinancing operations. Additional bonds eligible for regular purchases.
Euro Area	Dec. 3, 2015 13:45	Deposit Facility	−0.20	−10	0	18	Other rates unchanged. Continue asset purchase program until March 2017. Expected was an increase in monthly purchase volume.
Euro Area	Sep. 4, 2014 13:45	Deposit Facility	−0.10	−10	−10	−4	Interest rate on main refinancing operations decreased by 10 bps to 0.05%. Interest rate on marginal lending facility decreased by 10 bps to 0.30%. Introduction of asset-backed securities and covered bond purchase programs.
Euro Area	Jun. 5, 2014 13:45	Deposit Facility	−0.0	−10	0	−3	Interest rate on main refinancing oeprations decreased by 10 bps to 0.15%. Interest rate on marginal lending facility decreased by 35 bps to 0.40%. Introduction of targeted longer-term refinancing operations.

(*continued*)

Table 1. (Continued)

Country	Date, Time	Policy Rate Type	Prev. Pol. Rate (%)	Pol. Rate Change (bps)	Surpr. Comp.[a] (bps)	Gov. Yield Change[b]	Additional Policy Announcements
Japan	Jan. 29, 2016 04:38	Current Accounts	0.00	−10	−10	−7	Three-tier system, determining which rate applies. Loan Support Program, Funds-Supplying Operations.
Sweden	Feb. 11, 2016 09:30	Repo Rate	−0.35	−15	−5	−4	Continue government bond purchases for six months, reinvest yields from current portfolio.
Sweden	Jul. 2, 2015 09:30	Repo Rate	−0.25	−10	−10	−5	Extend government bond purchases by SEK 45 billion.
Sweden	Mar. 18, 2015 09:30	Repo Rate	−0.1	−15	−15	−11	Extend government bond purchases by SEK 30 billion.
Sweden	Feb. 12, 2015 09:30	Repo Rate	0.00	−10	−10	−15	Extend government bond purchases by SEK 10 billion.
Denmark	Feb. 5, 2015 16:00	Cert. Deposit	−0.50	−25	−25	−5	—
Denmark	Jan. 29, 2015 16:00	Cert. Deposit	−0.35	−15	−15	−1	—
Denmark	Jan. 22, 2015 16:00	Cert. Deposit	−0.20	−15	−15	−2	—
Denmark	Jan. 19, 2015 16:00	Cert. Deposit	−0.05	−15	−15	−6	—
Denmark	Sep. 4, 2014 16:00	Cert. Deposit	0.05	−10	0	−5	—
Denmark	Jul. 5, 2012 16:00	Cert. Deposit	0.05	−25	0	−7	—

[a]Daily % change of five-year government bond yields. Based on data from Thomson Reuters Datastream, using constant maturity series (e.g., "GVSD03(CM05)"). Euro-area yields are approximated with German government bond yields.
[b]The surprise component of the announced interest rate cut is computed as the difference between the announced cut and the median from the Bloomberg survey conducted among economists at domestic and international banks, usually on Fridays prior to the announcement.

announcing country through the expectations channel of the term structure, but we remove this direct impact by considering only how long-term interest rates respond in non-announcing countries. There may still be indirect effects through standard policy spillovers, however. These could be in the form of spillovers of policy rate moves to foreign policy rate moves, to expected future policy moves (i.e., lower for longer expectations), or through the term premium. Note first that our model suggests that if policy interest rates move in other countries in response to an unexpected interest rate cut below zero, the isolated effect on long-term yields would be smaller when it is closer to the lower bound, *given* the believed lower bound, which is opposite of the effect of a reduction in the believed lower bound. We briefly discuss the implications of other types of spillovers for our identification strategy here.

Consider the case where an unexpected negative interest rate announcement influences long-term yields in other countries through the expected duration of the other countries' own low-interest-rate regime. Even if an unexpected cut below zero does not lead to a higher likelihood of policy rate cuts in other countries, it could still lead market participants to believe that policy rates will stay low for longer than previously believed. This channel can be captured in our theoretical framework by assuming that foreign interest rate cuts reduce α and thus reduce the long-term average interest rate level. Alternatively, one could assume that the shock to the second-period interest rate ε_{t+1} has a non-zero mean μ, with foreign negative interest rate announcements lowering μ. These two formulations give identical results, so we focus here on the case where α changes. One finds that a decrease in α lowers long-term yields,

$$\frac{\partial R_t}{\partial \alpha} = \frac{1}{2}\left(1 - \Phi(c_{t+1})\right) > 0. \tag{12}$$

This effect is stronger if interest rates are higher:

$$\frac{\partial^2 R_t}{\partial \alpha \partial r_t} = \frac{1}{2\sigma}\phi(c_{t+1}) > 0. \tag{13}$$

Thus, just as in the case of an effect of foreign interest rate cuts on domestic short rates, we find that a potential "lower for longer" effect of foreign interest rate cuts would induce a positive correlation between changes in long-term yields and interest rate levels.

Now consider spillovers through term premiums. The term premium in the model of Ruge-Murcia (2016) plays no important role: shocks to the term premium are assumed to be uncorrelated with all other shocks in the model. In practice, however, movements in term premiums are important drivers of bond yield movements and likely to be endogenous, notably to policy moves and expectations. For example, announcements of cuts into negative territory were on some occasions accompanied by announcements of quantitative easing, and the associated portfolio balance effects on term premiums on government bonds could spill over to term premiums on assets in other countries. Moreover, it is often argued that cuts below zero are considered a sign of desperation in financial markets, resulting in higher risk premiums. This could increase the term premium but could also reduce it if the underlying instruments are considered safe (flight to safety). The empirical literature finds that term premiums are indeed an important source of variation and that they respond to monetary policy. For example, Crump, Eusepi, and Mönch (2016) find that for maturities greater than three years, term premiums explain more than 90 percent of the variation in the changes in bond yields. Abrahams et al. (2016) and Crump, Eusepi, and Mönch (2016) also find that term premiums increase in response to expansionary monetary policy shocks.

These findings are consistent with those of Rudebusch and Swanson (2012), who link an expansionary monetary policy surprise to expectations that economic activity will be higher and hence that marginal utility will be lower in the period when nominal bonds pay off. This makes bonds a less attractive investment, lowering their prices through a higher term premium. In contrast, Hanson and Stein (2015) argue that the decline in long-term real forward rates to surprise policy rate cuts is due to declines in the term premium—i.e., short-term rates and term premiums move in the same direction in response to monetary policy shocks. They build a model where the reaction of term premiums to policy shocks is due to the presence of yield-oriented investors. In response to a surprise interest rate cut, these investors rebalance their portfolios toward longer-term bonds in order to mitigate the decline in their average portfolio yield. The resulting increased demand for longer-term bonds raises their prices and lowers the term premium.

If monetary policy surprises affect term premiums, the relevance of this for identification depends on how it is related to the interest rate level. We are not aware of any studies that directly examine this link. Neither of the theoretical contributions mentioned above model the lower bound explicitly. Introducing a lower bound in these models could limit the response of term premiums to shocks, which would suggest that, if anything, the spillovers through term premiums could be either independent of the interest rate level or increasing in this level. This is opposite of the relation with the interest rate level suggested by changes in the believed lower bound. Potential spillovers through the term premium should therefore not affect our identification strategy.[7]

In general, the lower the interest rate level, the less standard policy spillovers should be able to influence long-term yields because of the lower-bound constraint. In contrast, the closer a country is to the lower bound, the larger the potential effect of a downward shift in the believed lower bound on interest rates. If an announcement has larger effects on long-term interest rates in countries that are closer to the lower bound than in countries that are further away from the lower bound, it implies that the believed lower bound is shifting down and that this effect is outweighing possible opposite effects of standard spillovers through policy expectations and risk premiums. We use these observations to formulate a sufficient condition for belief effects on long-term interest rates to be present.

To make this point more formally, consider a central bank that announces a surprise cut of its policy rate into negative territory, $dr_t < 0$. The effect on the long-term yield R_{it} in another country i is

$$\frac{dR_{it}}{dr_t} = \frac{\partial R_{it}}{\partial r_{it}} \frac{dr_{it}}{dr_t} + \frac{\partial R_{it}}{\partial \bar{r}_{it}} \frac{d\bar{r}_{it}}{dr_t}. \tag{14}$$

The first term on the right-hand side reflects a standard policy spillover effect through expected policy actions or the term premium.[8] The second term reflects the effect of changes in the believed

[7]Their presence could set the bar higher for identification, however.

[8]We have not included other standard spillovers such as changes in policy expectations (lower for longer) or changes in the term premium here. These effects

lower bound, and it is this effect that we are interested in identifying. If short-term interest rates do not respond to cuts in policy rates in other countries, then the identification problem is solved.[9] This would be a reasonable expectation if the announcing country is small in financial markets in comparison with country i.[10] It is nevertheless unlikely to be a good description of the data when the announcing country is influential in global financial markets or if the announcement is unusual and commands attention. We therefore rely on the predictions of the extended Ruge-Murcia (2006) model for identification. Consider how the above change in long-term yields depends on the interest rate level in country i. From (14), we have

$$\frac{d^2 R_{it}}{dr_t dr_{it}} = \frac{\partial^2 R_{it}}{\partial r_{it}^2}\frac{dr_{it}}{dr_t} + \frac{\partial R_{it}}{\partial r_{it}}\frac{d^2 r_{it}}{dr_t dr_{it}} + \frac{\partial^2 R_{it}}{\partial \bar{r}_{it}\partial r_{it}}\frac{d\bar{r}_{it}}{dr_t} + \frac{\partial R_{it}}{\partial \bar{r}_{it}}\frac{d^2 \bar{r}_{it}}{dr_t dr_{it}}. \tag{15}$$

The first two terms on the right-hand side capture traditional spillover effects on interest rates and how these vary with the interest rate level in country i. The last two terms capture how the effect of changes in the believed lower bound on long-term interest rates depends on the interest rate level. We now make assumption 2:

ASSUMPTION 2. *Short-term interest rates in country i either do not respond to changes in the policy rate of the announcing country or move in the same direction as policy rates in the announcing country:* $dr_{it}/dr_t \geq 0$.

Given business-cycle synchronization and the positive correlation of central bank policy rates across advanced economies, we consider

could easily be included, however, by modeling the effects on future policy expectations through a shift in the mean error of the interest rate process. A spillover through the term premium could be included through $d\theta_t/dr_t$.

[9]Note that in the version of the Ruge-Murcia (2006) model discussed in section 2, $dr_{it}/dr_t = 0$ holds by assumption. However, the model can be easily generalized by including foreign interest rates as an exogenous variable in the stochastic process for the shadow short rate, equation (2).

[10]This is unless smaller central banks had information about future ECB policy decisions in advance and lowered rates in anticipation of an upcoming ECB rate cut. However, we find such informational advantages to be unlikely in our sample.

it likely that this assumption is fulfilled in the data. Assumption 2 together with equation (8) above, which says that $\partial^2 R_{it}/\partial r_{it}^2 > 0$, imply that the first term on the right-hand side of (15) is non-negative. Note that the model prediction in equation (8) has been validated in time-series studies by Ruge-Murcia (2006) for Japan and by Grisse (2015) for the United States. Related evidence is provided by the finding by Swanson and Williams (2014) and others that bond yield responses to macroeconomic data surprises have declined as interest rates approached the lower bound. We therefore expect equation (8) to also hold in the cross-section of advanced economies. These studies include all spillover effects, including those through term premiums and expectations of future policy action.

The second term on the right-hand side of (15) is positive. We know this because equation (7) says that $\partial R_{it}/\partial r_{it} > 0$. Moreover, $d^2 r_{it}/dr_t dr_{it} = dr_{it}/dr_t \geq 0$ by assumption 2.

Now consider the third term on the right-hand side of (15). Equation (10) above says that $\partial^2 R_{it}/\partial \bar{r}_{it} dr_{it} < 0$. Therefore, if $d\bar{r}_{it}/r_t > 0$, the third term is negative.

Assumption 1 together with equation (9) above, which says that $\partial R_{it}/\partial \bar{r}_{it} > 0$, imply that the last term on the right-hand side of (15) is non-positive.

In short, under assumptions 1 and 2, the model shows that the effect of changes in the believed lower bound on long-term interest rates decreases in the level of interest rates prevailing in a given country. In contrast, the effect of changes in expected future short-term interest rates or term premiums on long-term interest rates increases in the level of interest rates prevailing in a given country. We can thus design our empirical strategy around a sufficient condition for the presence of effects of changes in the believed lower bound on long-term interest rates. Specifically, we assess whether the effect of unexpected interest rate cuts below zero on long-term interest rates in other countries is negatively related to prevailing interest rate levels.

It should be stressed that a negative relationship between the interest rate level and the effect of foreign interest rate cuts into negative territory on long-term interest rates is a sufficient but not necessary condition for the presence of belief effects. This is because we observe only the combined effects in the response of long-term yields to announcements. If standard spillovers are present and

Figure 2. Overview of Recent Policy Rate Cuts Below Zero

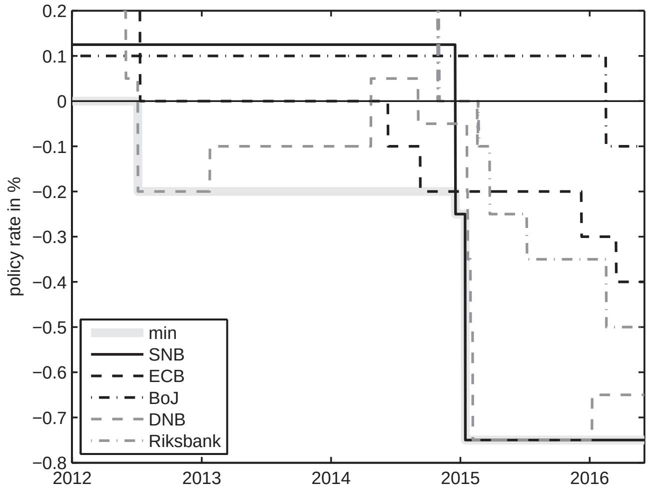

Notes: SNB: Three-month LIBOR target rate, Datastream series SWSNBTI. ECB: euro overnight deposit rate, Datastream series S20807. Bank of Japan: main policy rate, Datastream series Y74387. Danish National Bank: rate on certificates of deposit, Datastream series S838W7. Riksbank: policy rate, Datastream series SDRATEA.

positively related to the interest rate level, and if these outweigh the belief effects of the lower bound, the belief effects might still be there, but we would not see them. For the same reason, our approach does not allow us to identify the size of the believed lower-bound effects independently of other spillovers.

3.2 The Events

Table 1 lists the relevant policy rate cuts into negative territory undertaken so far, and figure 2 illustrates the time paths of the policy rates of the sample countries.

A downward shift in the believed lower bound and the traditional spillovers associated with these announcements should be accompanied by a decline in interest rates in non-announcing countries. As a first assessment of the data, we have plotted the movements of two-year yields, based on data on transactions-based intraday interest rate swap rates from Bloomberg, in the two and a half hours following each of the announcements that were unexpected.[11] Because of the global nature of the sample, some of these markets had very low liquidity and few transactions around the time of the announcements, and the data are hence not always very informative. A first look reveals that the effects of the announcements were highly heterogeneous across both countries and events. Below, we discuss the events in more detail and present those plots that best illustrate the findings of the model.[12]

The first policy rate cut into negative territory by the Danish central bank in 2012 came just hours after a surprise ECB rate cut in positive territory, and interest rates had therefore started to decline globally well before the Danish announcement. Table 1 shows that the cut was largely expected despite the novelty of moving policy rates below zero, perhaps also reflecting that the Danish central bank had been preparing a foundation for moving into negative territory. The subsequent ECB rate cut in June 2014 was also largely expected, according to the Bloomberg survey. The use of negative policy rates had been discussed, and the debate had evolved in the months preceding the actual cut; thus, possible effects on the believed lower bound would already have taken place by the time the ECB made its announcement (see also Lemke and Vladu 2016).

The following announcements were different. In September 2014, the ECB took the policy rate to the previous global low of –20 basis points that Denmark had already reached, and this cut was largely unexpected. As partly reflected in figure 3, this rate cut was associated with yield declines in the euro area as well as in the United States, the United Kingdom, Canada, Hungary, Poland, Switzerland, Denmark, Sweden, and Norway. Yields were declining for about one hour after the announcement (roughly until the beginning of the press conference). While these declines are consistent with a change in the market-perceived lower bound, it is common for global

[11]We focus on two-year yields on interest rate swaps; the picture is similar for five-, seven-, and ten-year yields.

[12]The plots for the remaining events are presented in the online appendix, in the section titled Appendix Figures.

Figure 3. Response of Two-Year Yields on Interest Rate Swaps around the ECB's Policy Announcement of September 4, 2014

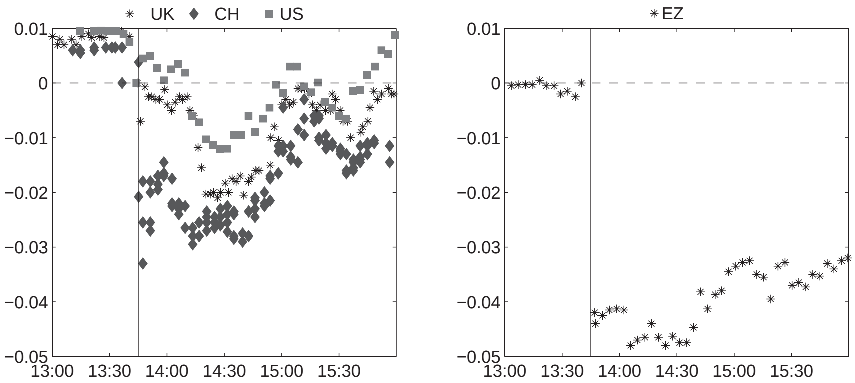

Notes: The cut was unexpected and included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

markets to react to ECB announcements, and a positive correlation of cross-country interest rate movements is not surprising.[13]

The SNB's 25 basis point rate cut to –0.25 percent in December 2014, after an unscheduled meeting, was unexpected. It resulted in a drop in Swiss long-term yields of about 8 basis points and was also associated with modest yield declines in the euro area, the United States, United Kingdom, Canada, Hungary, Poland, and Denmark, as partly shown in figure 4. Some yields subsequently increased by the end of the event window. Interestingly, this was mainly observed in countries that were away from the lower bound at that time. SNB announcements do not normally have large effects on foreign interest rates. The fact that this announcement was unusual, in that it reached a new global low of negative interest rates, may have made it different.

[13]At the press conference starting at 2:30 p.m. local time, the introduction of an asset purchase program was also announced. However, yields declined particularly after the release of the press statement (which mentioned only the interest rate decision) and *before* the press conference.

Figure 4. Response of Two-Year Yields on Interest Rate Swaps around the SNB's Policy Announcement of December 18, 2014

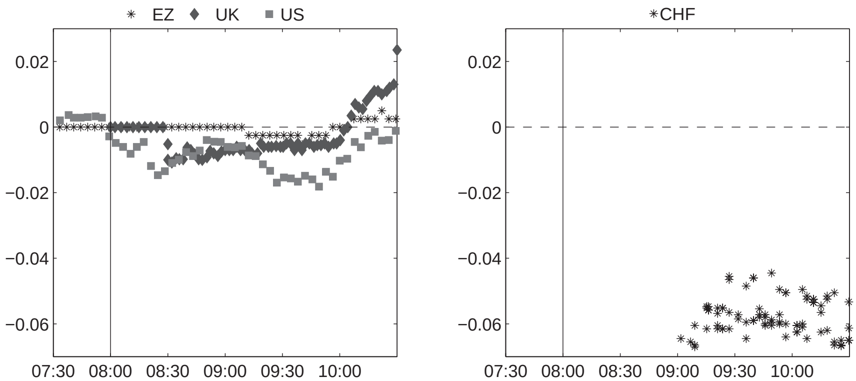

Notes: The cut was unexpected and included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

The SNB's 50 basis point rate cut to –0.75 percent in January 2015 was also fully unexpected and perhaps even more surprising, given how far this cut went beyond the previous global low. This cut was also associated with a decrease in Swiss long-term rates. The right panel in figure 5 shows that the two-year rate initially dropped about 20 basis points. The five-year yield dropped by 8 basis points the following day (see table 1). Yield declines were also observed in the euro area, the United States, United Kingdom, Canada, Denmark, and Norway, as shown in the left panel of figure 5. The announced interest rate cut accompanied the surprise announcement of the discontinuation of the 1.20 EUR/CHF exchange rate floor, which led to sharp appreciation in the Swiss franc and a global spike in FX volatility and risk-aversion measures. Therefore, while the observed yield movements are consistent with a change in the market-perceived lower bound, they are likely to reflect spillovers on term premiums and expectations that could have reinforced, or attenuated, the effect on longer rates in other

Figure 5. Response of Two-Year Yields on Interest Rate Swaps around the SNB's Policy Announcement of January 15, 2015

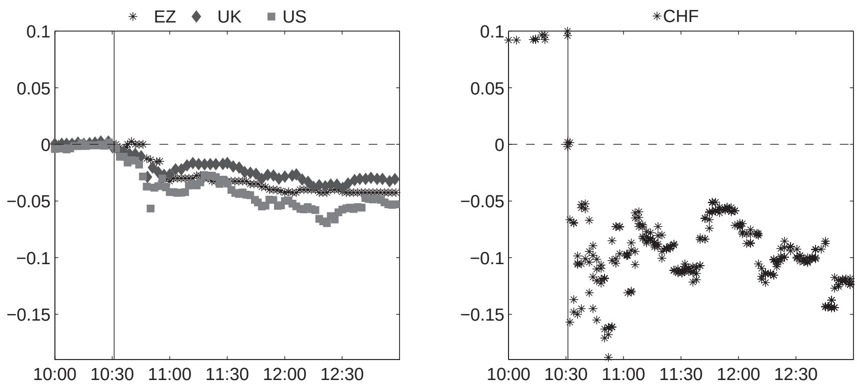

Notes: The cut was unexpected and included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

countries, underlining the need for our more rigorous approach to identification.[14]

Just days after the Swiss announcement, the Danish central bank cut its policy rate and continued to cut it in three further separate steps over the subsequent month, reaching a negative policy rate level equal to the global low of the Swiss policy rate by February 5, 2015. Each of these steps was unexpected according to the Bloomberg survey. In general, yields dropped in Denmark, but the responses in other countries were mixed, and we do not depict them here.[15]

Sweden's Riksbank subsequently cut its policy rate on February 12, March 18, and July 2, 2015, reaching a policy rate of –35

[14]Yields in Hungary and Poland increased, potentially reflecting a perceived deterioration in financial conditions in these countries, where the private sector had borrowed heavily in Swiss francs. Because of the Swiss franc appreciation, the local currency value of these debts rose.

[15]The figures are available in the online appendix.

basis points. These announcements were also unexpected, and they did not reach new global lows. The announcements were associated with important drops in Swedish yields and some downward movement in rates elsewhere. It is unlikely that the Danish and Swedish announcements in 2015 influenced the believed lower bound as much as the preceding Swiss announcements to which they were reacting, however, because they were less novel.

The ECB cut its leading policy rate further below zero on December 3, 2015, but this cut was fully expected according to the Bloomberg survey. On January 29, 2016, the Bank of Japan cut its leading policy interest rate to 10 basis points below zero. This was fully unexpected. Leading voices at the Bank of Japan had communicated just days earlier that negative interest rates were out of the question in Japan. Financial markets in Japan seemed completely unprepared, which triggered an increase in risk premiums and risk-off. While the cut was quantitatively small, the immediate impact on the believed lower bound in Japan could have been large because of the change in attitude inside the Bank of Japan that this cut could have been taken to signal. The cut was indeed associated with important long-term yield declines in Japan. Figure 6 shows that in the hours following the announcement, the two-year yield declined by 8 basis points. In the days that followed, long yields declined even more. Subsequent Bank of Japan communications have partly explained the unexpectedly large declines in long-term yields with downward shifts in the believed lower bound.[16] The announcement by the Bank of Japan was also associated with yield declines in the United States, the euro area, Australia, and New Zealand, but these declines were modest. As rates had already gone much further into negative territory elsewhere, revisions to the believed lower bound in other countries may have been small.

As in the case for the ECB, surprise announcements by the Bank of Japan normally affect markets globally, even if the believed lower bound is unaffected, calling for a more rigorous analysis of the data.

Sweden's Riksbank again cut its policy rate on February 11, 2016, but the unexpected component of this cut was only 5 basis points.

[16] See Kuroda (2016).

Figure 6. Response of Two-Year Yields on Interest Rate Swaps around the Bank of Japan's Policy Announcement of January 29, 2016

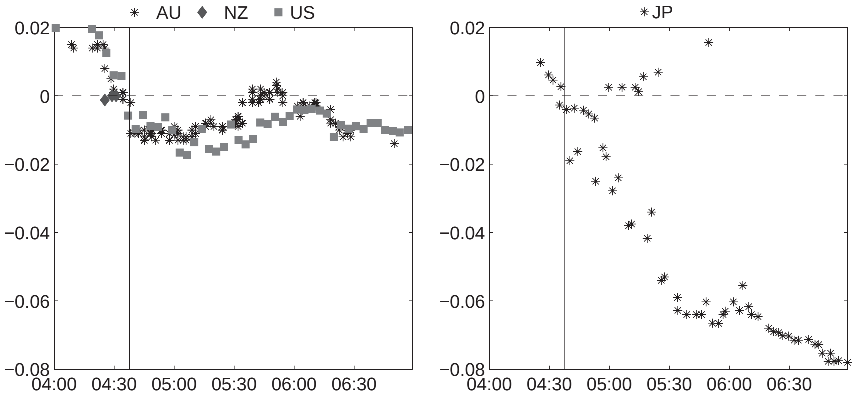

Notes: The cut was unexpected and included in our sample. Vertical lines mark the time of the announcement. Yields are normalized by adding a constant so that the average of the last two observations before the announcement is zero. Yields are based on transaction prices from Bloomberg.

Yields in other countries tended to moderately decline after this announcement.[17]

Finally, the ECB cut its leading rate by 10 basis points to –40 basis points on March 10, 2016. This cut was expected, and we do not learn much from the market reaction to it. Yields, if anything, tended to increase in response.[18] Interestingly, during the press conference following the announcement, the ECB President answered a question about how low interest rates could go with remarks about adverse consequences for the banking system of negative interest rates and the desire of the ECB's Governing Council not to signal that rates could go as "low as we want."[19] At about this time, yields in other countries moderately increased.

[17] Not shown, but see figure 4 in the online appendix.

[18] The ECB did not introduce a tiered system for the interest rate on central bank reserves on deposit with the ECB, as had been expected, so the announcement as a whole may have been less of a stimulus than expected.

[19] Transcript from the ECB press conference on March 10, 2016, available at http://www.ecb.europa.eu/press/pressconf/2016/html/is160310.en.html.

Overall, global market movements around the announcements suggest a tendency for foreign yields to decline when announced cuts were unexpected. The data also suggest that there is a lot of variation in the market response across these unexpected announcements and across countries that requires further investigation.

3.3 Panel Investigation

We now turn to a more formal panel data analysis of the lower-bound location channel.

3.3.1 Specification

Our baseline panel specification is as follows:

$$\Delta R_{i,t} = \beta_0 + \beta_1 \Delta r_t + \beta_2 r_{i,t-1} + \beta_3 \Delta r_t \times r_{i,t-1} + \varepsilon_{i,t} \tag{16}$$

$$\Delta R_{i,t} = \beta_0 + \beta_1 \Delta r_t + \beta_2 R_{i,t-1} + \beta_3 \Delta r_t \times R_{i,t-1} + \varepsilon_{i,t}. \tag{17}$$

The index t runs across the selected twelve announcement dates; the index i runs across the countries in our sample excluding the country announcing at time t. This implies that the sample composition changes across announcements. The surprise component of the change in the policy rate in the announcing country, Δr_t, proxies for the revision of market views about the lower bound but also captures the standard policy spillovers through interest rates discussed above. The distance to the lower bound in the home country is measured in two alternative ways. We first use the level of the home-country policy rate $r_{i,t-1}$ in specification (16). Alternatively, we use the long-term yield $R_{i,t-1}$ in specification (17) to capture the proximity of the lower bound in country i over the lifetime of the long-term bond. Note that the coefficient of proportionality γ from equation (12) is subsumed within the coefficients β_1 and β_3, while a potential non-zero mean of $\upsilon_{i,t}$ would be subsumed within β_2.

Focusing on specification (16) for a moment, the effect of Δr_t on $\Delta R_{i,t}$ is given by $\beta_1 + \beta_3 r_{i,t-1}$. The hypothesis that a downward revision in market views about the lower bound is associated with a decline in long-term yields and that, from equation (10), this effect is stronger for countries where interest rates are more constrained then corresponds to $\beta_3 < 0$: if $\beta_3 < 0$, then a lower value of $r_{i,t-1}$

is associated with an increase in $\partial R_t/\partial \bar{r}$. Following the discussion in section 3.1, if $\beta_3 < 0$, we can interpret this as evidence that the change in the believed lower bound affected long-term interest rates and that this lower-bound effect has outweighed any possible spillover effect through short rate changes.

The coefficient β_1 represents $\partial R_t/\partial \bar{r}$ under the condition that $r_{i,t-1} = 0$. Note that we cannot include episode-specific fixed effects because of the presence of Δr_t in the regression. Therefore, we cannot account for global factors that affect the yields of all countries in t in the same way—for example, changes in risk sentiment or globally relevant economic news. The coefficient β_1 is then difficult to interpret because it reflects both the effect of negative interest rate announcements when $r_{i,t-1} = 0$ and other episode-specific global movements in yields. A priori, one would not expect the change in yields in any particular episode to depend on the initial yield level; thus, we expect $\beta_2 = 0$.[20]

3.3.2 Sample and Data

Our sample of countries comprises Australia, Canada, the Czech Republic, Denmark, the euro area, Hungary, Japan, New Zealand, Norway, Poland, Sweden, Switzerland, the United Kingdom, and the United States. This sample is chosen to include countries where policy rates have been likely to be directly constrained by the believed location of the lower bound (Denmark, Switzerland, Sweden, Japan, and the euro area), countries that have been somewhat constrained (the United Kingdom and United States), and some that have been more comfortably away from the lower bound, such as Australia and New Zealand, as well as Poland and Hungary, to some degree.[21] We have a total of 155 observations, based on thirteen countries (fourteen less the announcing country) times twelve announcements,

[20]However, the effect of the (policy) interest rate level $r_{i,t-1}$ on the yield change is $\beta_2 + \beta_3\Delta r_t$. Because we are focusing on negative interest rate announcements, we have $\Delta r_t < 0$ by construction. Therefore, the coefficient β_2 in itself (the effect of $r_{i,t-1}$ given that $\Delta r_t = 0$) is not very meaningful. If $\beta_3 < 0$ and if the effect of $r_{i,t-1}$ is small across countries and across announcements, then we expect $\beta_2 < 0$.

[21]We do not include euro-area countries individually since the lower bound on the short-term interest rate refers to monetary policy decisions and is hence the same across the euro area.

excluding Denmark for ECB's September 2014 announcement, where Denmark subsequently cut its own policy rate.

As the dependent variable, we use the daily change in yields on interest rate swaps ICAP series from Datastream, which according to Datastream are computed using data as of 5:00 p.m. CET for all sample countries. The short rate in the swap is the three-month rate. The advantage of this data set is that the yield is computed at exactly the same time across the sample countries, allowing us to capture the announcements. We use daily data and a one-day event window because there is very little liquidity around our negative interest rate events and because the one-day window allows yields in these countries to respond when markets open. Daily data are available for a larger set of countries. Depending on the announcement time, the one-day window is split into a short (a few hours) pre-announcement window and a longer post-announcement window (for the Bank of Japan announcement) or into a short (a few hours) post-announcement window and a longer pre-announcement window (for those European announcements that are made in the afternoon).

As explanatory variables, we use the surprise component of interest rate announcements (Δr_t), as computed from the Bloomberg Survey and listed in table 1. As measures of domestic interest rate levels, we use the central bank policy rate ($r_{i,t-1}$) or the interest rate on the swap contract that is also used to compute the dependent variable ($R_{i,t-1}$), as of the day prior to the announcement. The latter captures the relevance of the lower bound over the lifetime of the swap contract. Summary statistics for the regression variables are presented in table 2.[22]

3.3.3 Results

Table 3 contains the baseline regression results. The dependent variable is the change in the long-term interest rate in basis points. The first four columns contain the results for specification (16) when the prevailing policy rate is used as a measure of the believed distance to the lower bound for all four different maturities of the dependent variable that we consider. The last four columns display the results

[22]Data sources and definitions are summarized in the online appendix, table 1.

Table 2. Summary Statistics

	Mean	**St. Dev.**	**Min.**	**Max.**	**Obs.**
Monetary Policy Surprises, Δr_t (bps)	−17.92	12.15	−50.00	−5.00	12
Central Bank Policy Rates, $r_{i,t-1}$(%)	0.89	1.14	−0.75	3.50	155
Yield Change, $\Delta R_{i,t}$ (bps)					
Two-Year Yields	1.06	1.07	−0.76	4.17	155
Five-Year Yields	1.37	1.09	−0.56	4.49	155
Seven-Year Yields	1.57	1.09	−0.34	4.61	155
Ten-Year Yields	1.83	1.07	−0.07	4.73	155
Yield Level, $R_{i,t-1}$ (%)					
Two-Year Yields	−1.56	2.97	−9.11	8.10	155
Five-Year Yields	−1.80	3.87	−12.01	11.52	155
Seven-Year Yields	−1.82	4.25	−14.25	11.45	155
Ten-Year Yields	−1.78	4.72	−15.80	11.60	155

Notes: We include only days of negative interest rate announcements and only dates with non-zero announcement surprises based on the median of the Bloomberg survey. We pool the data across these twelve announcements and fourteen countries, but for each announcement we exclude the announcing country, leaving thirteen countries per announcement. Interest rate levels are measured in percentage; interest rate changes are measured in basis points. Yields refer to daily data as of 5:00 p.m. CET on ICAP yields on interest rate swaps with a three-month variable leg, obtained from Thomson Reuters Datastream.

for specification (17) where the long-term rate that is also used to compute the dependent variable is used as a measure of the believed distance to the lower bound over the lifetime of the fixed-income investment in question.

The constant term is negative, suggesting that, on average, long-term interest rates dropped on announcement dates by 2 to 3 basis points depending on the maturity and specification. In contrast, we find that the direct effect of unexpected rate cuts is negatively proportional to the interest rate cut—often significantly so. Combined with the constant term, this could reflect average changes in macroeconomic conditions to which the rate cut in the announcing country

Table 3. Panel Regression Results: Baseline

	Specification (16)				Specification (17)			
Maturities	**2**	**5**	**7**	**10**	**2**	**5**	**7**	**10**
Δr_t	−0.01	−0.06	−0.08	−0.10	−0.01	−0.05	−0.07	−0.08
	(0.04)	(0.07)	(0.08)	(0.09)	(0.04)	(0.07)	(0.08)	(0.08)
$r_{i,t-1}$	−1.09***	−1.00**	−0.98**	−0.89**				
	(0.39)	(0.46)	(0.43)	(0.42)				
$\Delta r_t \times r_{i,t-1}$	−0.05***	−0.04**	−0.03**	−0.03*				
	(0.01)	(0.02)	(0.01)	(0.01)				
$R_{i,t-1}$					−0.98**	−0.84	−0.88*	−0.88*
					(0.45)	(0.53)	(0.48)	(0.46)
$\Delta r_t \times R_{i,t-1}$					−0.04***	−0.03**	−0.03***	−0.03**
					(0.01)	(0.01)	(0.01)	(0.01)
Constant	−1.50**	−2.54**	−2.92**	−3.19*	−1.40*	−2.25*	−2.37	−2.30
	(0.73)	(1.20)	(1.40)	(1.65)	(0.75)	(1.25)	(1.47)	(1.73)
No. Obs.	155	155	155	155	155	155	155	155
R^2	0.08	0.09	0.10	0.10	0.07	0.08	0.09	0.09

Notes: Cluster-robust standard errors are in parentheses. Statistical significance is indicated by *** if $p < 0.01$, ** if $p < 0.05$, and * if $p < 0.1$. $\Delta R_{i,t}$ for the announcing country has been excluded.

is reacting, but it could also reflect some average effect of the cut that is independent of the size of the cut or home-country interest rate level.

The estimated coefficients on the interaction term $\Delta r_t \times r_{i,t-1}$ have the expected negative sign and are highly statistically significant across maturities and specifications. The further away from the lower bound a country is when an unexpected policy rate cut is announced in another country, the less the home-country long-term interest rate decreases on announcement dates (or the more the long-term interest rate in the home country increases on such dates). This is consistent with the theoretical prediction that negative interest rate announcements have affected yields globally by affecting market views about the lower bound and that the strength of this effect depends on the proximity of the lower bound prior to the announcements.

Quantitatively, the composite average effect of the announced cuts is not large, but as already stressed in section 3.1, the total predicted effect does not contain information about the net effect of changes in the believed lower bound. It only suggests that it is there, as it must be outweighing other effects that pull in the opposite direction. Moreover, it is averaged over a set of announcements that is likely to contain episodes where the believed lower bound was not affected (attenuation bias). The isolated effect of changes in the believed lower bound is likely to be substantially larger, but we cannot assess by how much.

Overall, the baseline results lend empirical support to the presence of lower-bound belief effects on average for the series of announcements of unexpected cuts below zero.

3.4 Robustness

We conducted a number of robustness checks for the panel approach. First, we ran the regressions while including only those announcements where new global lows were reached for policy rates: since the Danish announcement in 2012 was expected and hence cannot capture changes in the believed lower bound, this leaves the two Swiss announcements of December 2014 and January 2015. The results are shown in table 4. The results for the parameter estimates of β_3 are negative and larger in size than those of the baseline specification,

Table 4. Panel Regression Results: New Global Lows

	Specification (16)				Specification (17)			
Maturities	**2**	**5**	**7**	**10**	**2**	**5**	**7**	**10**
Δr_t	0.18***	0.26***	0.27***	0.28***	0.18**	0.25**	0.27**	0.30*
	(0.07)	(0.09)	(0.08)	(0.09)	(0.08)	(0.12)	(0.13)	(0.15)
$r_{i,t-1}$	−2.78	−4.13*	−4.07*	−4.40*				
	(1.75)	(2.34)	(2.19)	(2.31)				
$\Delta r_t \times r_{i,t-1}$	−0.08*	−0.11*	−0.10*	−0.10*				
	(0.04)	(0.06)	(0.06)	(0.06)				
$R_{i,t-1}$					−2.20	−2.55	−2.22	−2.44
					(1.98)	(2.53)	(2.49)	(2.68)
$\Delta r_t \times R_{i,t-1}$					−0.07	−0.07	−0.06	−0.06
					(0.05)	(0.07)	(0.06)	(0.07)
Constant	6.35**	10.75***	12.10***	13.06***	6.21*	10.47**	11.80**	13.49**
	(2.66)	(3.39)	(3.32)	(3.50)	(3.21)	(4.72)	(5.05)	(6.00)
No. Obs.	26	26	26	26	26	26	26	26
R^2	0.26	0.29	0.33	0.33	0.21	0.22	0.25	0.25

Notes: Only "new global low" surprise announcements are included (Switzerland, December 2014 and January 2015). Because conventional standard errors for these regressions are lower than cluster-robust standard errors (for all coefficients), we here report conventional standard errors in parentheses. Statistical significance is indicated by *** if $p < 0.01$, ** if $p < 0.05$, and * if $p < 0.1$. $\Delta R_{i,t}$ for the announcing country has been excluded.

as could be expected if these announcements had stronger effects on the believed lower bound than the other announcements included in the sample. However, the level of significance drops considerably. This is not surprising given the substantial reduction in observations and hence degrees of freedom.[23] In contrast, the constant term in this regression turns positive and becomes large, perhaps reflecting generally increasing risk premiums around those dates in non-Swiss long-term yields. The signs of β_3 are consistent with the hypothesis that the two SNB rate cuts below zero changed the believed lower bound abroad, however.

Second, while a negative interest rate announcement in another country may influence average perceptions about how much further down the domestic central bank is willing to cut rates in the future, the influence on the believed lower bound could be smaller in countries where the policy rate was already lower than the rate in the announcing country. Excluding all observations for countries where the policy rate was already lower than the level reached by the rate cut in the announcing country, however, does not change the size and significance of the coefficients on the interaction terms.[24]

Third, one could worry that the direct transmission through short rates might be stronger for low-interest-rate countries because announcements of negative rate cuts were mostly in Europe, where rates were generally lower than in other parts of the world. Australia and New Zealand have the highest policy rates in our sample of countries. Both are above 2 percent (New Zealand as high as 3.5 percent) for all observations. Poland and Hungary also have high policy rates (close to 2 percent). These countries are included to add observations and variation in the interest rate level. It could, however, be that geography matters for monetary policy spillovers and that these countries saw less of a downward movement in their rates as a response to the rate announcements that we study simply because they are further away. To address this concern, we ran the regressions while excluding Australia and New Zealand and then

[23] In the specification including only the two SNB announcements, robust standard errors are markedly lower than conventional standard errors, indicating the presence of small-sample bias. In table 4, we therefore report conventional standard errors instead.

[24] The regression results with this exclusion are not shown here but are available in the online appendix, table 2.

also excluding Poland and Hungary (not shown).[25] Even though these exclusions strongly reduce the variation in the interest rate levels, the results are largely robust.[26]

Finally, we measured the distance to the believed lower bound by the level of the interest rate in our regressions, implicitly proxying the lower bound by zero, as we do not observe this level. The level of the lower bound is assumed to be shifting throughout the sample, however, and omitting this variation may create omitted-variable bias if shifts in the level of the believed lower bound are correlated with the error terms. There is no a priori reason why such a correlation should be expected. Moreover, when proxying the level of the believed lower bound with the level of the globally lowest interest rate, the size and significance of the coefficients on the interaction terms remain unchanged.[27]

4. Conclusion

This paper studies the effects of a change in the believed location of the lower bound on the transmission of short-term interest rates to longer-term yields. Using a simple theoretical term structure model originally proposed by Ruge-Murcia (2006), we show that if interest rate cuts below zero lead market participants to adjust downward their expectations regarding how far below negative interest rates can go, this can further reduce long-term interest rates over and above the impact of the policy rate cut itself.

Using an event study of yield reactions to selected negative interest rate announcements in a panel of advanced economies, we find evidence consistent with the hypothesis that long rates indeed have tended to drop in response to downward revisions in the market's believed location of the lower bound.

These findings have important implications both for monetary policy with low interest rates and for the use of negative interest rates as a monetary policy tool. The transmission of rate cuts into

[25] The results are not shown here but are available in the online appendix, tables 3 and 4.

[26] Only in the yield-level specification excluding all four countries does the interaction term lose significance. Yet, even here, β_3 is negative and large, just with large standard errors.

[27] Not shown, but the results are presented in the online appendix, table 5.

negative territory is enhanced if a rate cut lowers the believed location of the lower bound. Moreover, to the extent that central banks have or are believed to have private information about the location of the lower bound, communication that informs market participants about this location can have important effects on the monetary policy stance.

References

Abrahams, M., T. Adrian, R. K. Crump, E. Mönch, and R. Yu. 2016. “Decomposing Real and Nominal Yield Curves.” *Journal of Monetary Economics* 84 (December): 182–200.

Brunnermeier, M. K., and Y. Koby. 2016. “The Reversal Interest Rate: An Effective Lower Bound on Monetary Policy.” Unnpublished Manuscript.

Crump, R. K., S. Eusepi, and E. Mönch. 2016. “The Term Structure of Expectations and Bond Yields.” Staff Report No. 775, Federal Reserve Bank of New York.

Danthine, J.-P. 2016. “The Interest Rate Unbound?” Report, Brookings Institution.

Grisse, C. 2015. “The Zero Lower Bound and Movements in the Term Structure of Interest Rates.” *Economics Letters* 131 (June): 66–69.

Hanson, S. G., and J. C. Stein. 2015. “Monetary Policy and Long-Term Real Rates.” *Journal of Financial Economics* 115 (3): 429–48.

Kortela, T. 2016. “A Shadow Rate Model with Time-Varying Lower Bound of Interest Rates.” Research Discussion Paper No. 19, Bank of Finland.

Krugman, P. R. 1991. “Target Zones and Exchange Rate Dynamics.” *Quarterly Journal of Economics* 106 (3): 669–82.

Kuroda, H. 2016. “Re-anchoring Inflation Expectations via ‘Quantitative and Qualitative Monetary Easing with a Negative Interest Rate’.” Remarks at the Economic Policy Symposium held by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming, August 25–27.

Lemke, W., and A. L. Vladu. 2016. “Below the Zero Lower Bound: A Shadow-Rate Term Structure Model for the Euro Area.” Discussion Paper No. 32/2016, Bundesbank.

Rudebusch, G. D., and E. T. Swanson. 2012. "The Bond Premium in a DSGE Model with Long-Run Real and Nominal Risks." *American Economic Journal: Macroeconomics* 4 (1): 105–43.

Ruge-Murcia, F. J. 2006. "The Expectations Hypothesis of the Term Structure when Interest Rates Are Close to Zero." *Journal of Monetary Economics* 53 (7): 1409–24.

Swanson, E. T., and J. C. Williams. 2014. "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates." *American Economic Review* 104 (10): 3154–85.

Discussion of "Lower-Bound Beliefs and Long-Term Interest Rates"

James D. Hamilton
University of California at San Diego

1. Introduction

Grisse, Krogstrup, and Schumacher (this issue) provide one of the first systematic evaluations of the effects of the negative interest rate policies recently adopted across various countries. The authors have made a welcome and useful contribution.

The authors use the Ruge-Murcia (2006) shadow-rate model as a theoretical framework for interpreting the evidence. I would first like to suggest a much more general way this could be done and then offer some comments on using event studies for empirical evaluation.

2. Shadow-Rate Models of the Lower Bound

Following the authors' notation, let r_t denote the short-term nominal interest rate, $\bar{r}$ the effective lower bound, and r_t^* the shadow rate. The latter is a theoretical construct that is allowed to be an arbitrarily large negative number. The idea is that when r_t^* is below $\bar{r}$, the observed short rate will be in the vicinity of $\bar{r}$:[1]

$$r_t = \max\{r_t^*, \bar{r}\}.$$

The authors employ Ruge-Murcia's (2006) description of the process followed by the shadow rate,

$$r_{t+1}^* = \alpha + \psi r_t + \varepsilon_{t+1}, \tag{1}$$

where $\varepsilon_{t+1} \sim N(0, \sigma^2)$. One problem with this specification is that it cannot explain a persistent episode with rates at the lower bound.

[1] Empirical models typically also incorporate the possibility of measurement error to allow some fluctuation of the observed rate around the constant $\bar{r}$.

Suppose for example that we're currently at a zero lower bound, with $r_t = \bar{r} = 0$. Then (1) implies that $E_t(r^*_{t+1}) = \alpha > 0$ and $E_t(r^*_{t+n}) > 0$ for all n. There is no way such a process could be consistent with interest rates that were observed to stay near zero for years.

A more popular specification in the literature (e.g., Krippner 2013, 2015; Christensen and Rudebusch 2015; and Wu and Xia 2016) takes the form

$$r^*_{t+1} = \alpha + \psi r^*_t + \varepsilon_{t+1}, \tag{2}$$

which implies $E_t(r^*_{t+n}) = \frac{\alpha(1-\psi^n)}{(1-\psi)} + \psi^n r^*_t$. If ψ is near unity and r^*_t is far below zero, then $E_t(r^*_{t+n})$ could remain negative for large n. News (in the form of realizations of ε) may change the expected date of "liftoff" above the lower bound $\bar{r}$ because it changes $E_t(r^*_{t+n})$.

To arrive at a general description of the term structure of interest rates for such a process, it's helpful to define a few additional terms. Let y_{nt} denote the continuously compounded yield on a zero-coupon bond maturing at $t+n$ (so that y_{1t} is just another symbol for r_t, the one-period rate). Let f_{nt} denote the n-period-ahead forward rate. This is a rate we could lock in at date t by selling an n-period bond and simultaneously buying an $(n+1)$-period bond:

$$f_{nt} = (n+1)y_{n+1,t} - ny_{nt}. \tag{3}$$

Note that from this definition, an n-period interest rate can be viewed as the average of the corresponding set of forward rates:

$$y_{nt} = n^{-1}(f_{n-1,t} + f_{n-2,t} + \cdots + f_{1t} + y_{1t}). \tag{4}$$

Note that (4) is true by the definition of a forward rate in (3), and does not make any assumptions whatever about investors' objectives or beliefs.

To calculate the predicted behavior of interest rates under a simple model of investors' preferences, consider first an economy that is currently far away from the lower bound, so that next period's short rate should equal the shadow rate: $r_{t+1} = r^*_{t+1}$. We could then buy a two-period bond at date t for price $\exp(-2y_{2t})$ and sell it at $t+1$ for $\exp(-r^*_{t+1})$ for a gross expected return of

$$E_t[\exp(2y_{2t})\exp(-r^*_{t+1})] = \exp(2y_{2t} - E_t r^*_{t+1} + \sigma^2/2),$$

where the term $\sigma^2/2$ is a consequence of Jensen's inequality.[2] The expected excess return from holding a two-period bond over a one-period bond is thus

$$2y_{2t} - E_t r^*_{t+1} + \sigma^2/2 - y_{1t} = f_{1t} - E_t r^*_{t+1} + \sigma^2/2,$$

with the last equality following from (3). We could define the expected excess return to be the term premium. Suppose that investors' tolerance for risk is correlated with the current shadow rate r^*_t and that the term premium could be written as $\lambda_0\sigma + \lambda_1\sigma r^*_t$ for some constants λ_0 and λ_1:[3]

$$f_{1t} - E_t r^*_{t+1} + \sigma^2/2 = \lambda_0\sigma + \lambda_1\sigma r^*_t. \tag{5}$$

Note that this specification includes the expectations hypothesis of the term structure as a special case when $\lambda_0 = \lambda_1 = 0$. The equilibrium condition (5) can alternatively be written

$$\begin{aligned} f_{1t} &= \alpha + \psi r^*_t - (\sigma^2/2) + \lambda_0\sigma + \lambda_1\sigma r^*_t \\ &= \alpha^Q + \psi^Q r^*_t - (\sigma^2/2), \end{aligned}$$

where $a^Q = \alpha + \lambda_0\sigma$ and $\psi^Q = \psi + \lambda_1\sigma$. In other words, investors could be viewed as if they take expectations of future r^*_{t+1} not using the objective process (2) but instead using the Q measure $r^*_{t+1} \sim N(\alpha^Q + \psi^Q r^*_t, \sigma^2)$:

$$f_{1t} = E^Q_t r^*_{t+1} - (\sigma^2/2).$$

We can likewise calculate a risk-adjusted expectation of the shadow rate n periods ahead:

$$E^Q_t(r^*_{t+n}) = \frac{\alpha^Q[1-(\psi^Q)^n]}{(1-\psi^Q)} + (\psi^Q)^n r^*_t. \tag{6}$$

If we assume that investors care about risk-adjusted returns as calculated by the Q-measure parameterization (6), then Wu and

[2]Since $-r^*_{t+1}|r^*_t \sim N(-E_t r^*_{t+1}, \sigma^2)$, $E_t \exp(-r^*_{t+1}) = \exp(-E_t r^*_{t+1} + \sigma^2/2)$.

[3]Section 2.1 in Hamilton and Wu (2014) illustrates how such a functional dependence could arise.

Figure 1. The Equilibrium Forward Rate under the Shadow-Rate Model

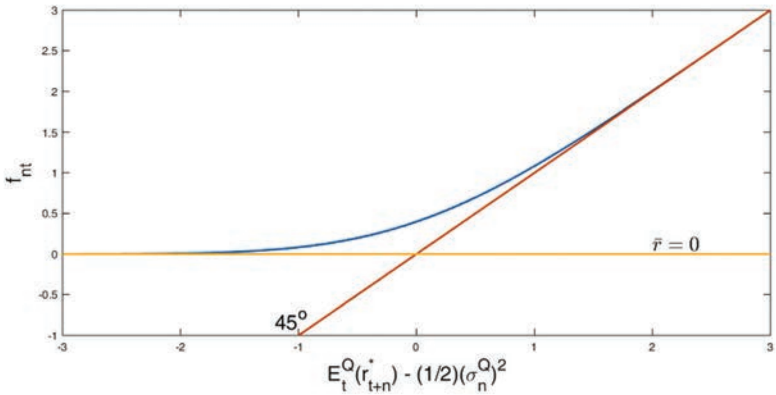

Notes: Horizontal axis: Q-measure expectation of the shadow rate n periods in the future (with σ_n^Q normalized at 1). Vertical axis: n-period-ahead forward rate.

Xia (2016) demonstrated that in equilibrium, forward rates can be approximated as

$$f_{nt} \simeq \bar{r} + \sigma_n^Q g\left(\frac{E_t^Q(r^*_{t+n}) - (1/2)(\sigma_n^Q)^2 - \bar{r}}{\sigma_n^Q}\right). \tag{7}$$

Here $g(z) = z\Phi(z) + \phi(z)$ for $\Phi(z)$ the cumulative distribution function for a standard normal variable and $\phi(z)$ the density, and

$$\begin{aligned}(\sigma_n^Q)^2 &= E_t^Q[r^*_{t+n} - E_t^Q(r^*_{t+n})]^2 \\ &= \frac{\sigma^2[1-(\psi^Q)^{2n}]}{1-(\psi^Q)^2}.\end{aligned}$$

To understand the intuition behind (7), suppose first that $g(.)$ was the identity function ($g(z) = z$). In this case (7) would simplify to $f_{nt} = E_t^Q(r^*_{t+n}) - (1/2)(\sigma_n^Q)^2$, corresponding to the case derived above when we are far from the lower bound. Figure 1 plots equation (7) for the general case when $g(z) = z\Phi(z) + \phi(z)$. When $E_t^Q(r^*_{t+n})$ is very far above the lower bound, $g(z)$ approaches the 45-degree line, and the forward rate is essentially the same as that

Figure 2. Effect of Changing the Lower Bound on the Forward Rate

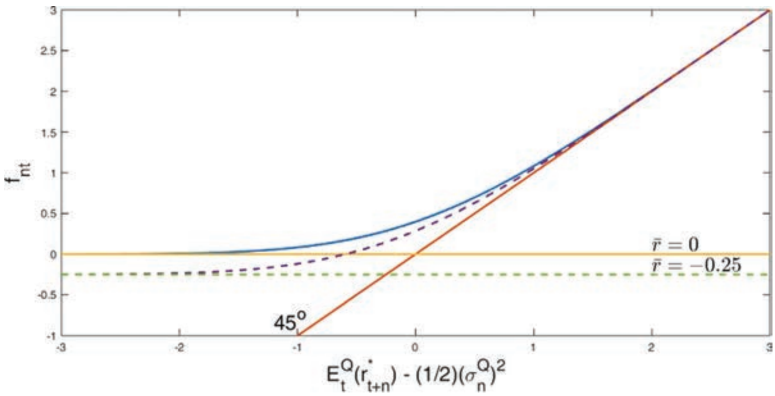

predicted in the absence of a lower bound. When $E_t^Q(r^*_{t+n})$ is far below the lower bound, the value for f_{nt} is essentially the lower bound $\bar{r}$ itself. The $g(.)$ function thus provides a smooth pasting to generate a forward rate that is larger than $\bar{r}$ for all n and asymptotes to $E_t^Q(r^*_{t+n}) - (1/2)(\sigma_n^Q)^2$.

We can also see immediately from (7) the effects of a change in the lower bound:

$$\frac{\partial f_{nt}}{\partial \bar{r}} = 1 - g'.$$

Here g' denotes the derivative of the function $g(.)$, which analytically turns out to be given by $g'(z) = \Phi(z)$ and is bounded between 0 and 1 for all z. Thus

$$\frac{\partial f_{nt}}{\partial \bar{r}} = 1 - \Phi\left(\frac{E_t^Q(r^*_{t+n}) - (1/2)(\sigma_n^Q)^2 - \bar{r}}{\sigma_n^Q}\right).$$

As $E_t^Q(r^*_{t+n}) \to -\infty$, $\partial f_{nt}/\partial \bar{r} \to 1$, whereas when $E_t^Q(r^*_{t+n}) \to \infty$, $\partial f_{nt}/\partial \bar{r} \to 0$. The effect of changing $\bar{r}$ on the forward rate is illustrated in figure 2. Recalling (4), a decrease in the lower bound $\bar{r}$ should lower short-term yields nearly one-for-one but have a much more modest effect on long-term yields.

I have followed Grisse, Krogstrup, and Schumacher (this issue) up to this point in assuming that the shadow rate was described

Table 1. Changes in German Term Structure Associated with Changes in the ECB Lower Bound

Date	Old Rate	New Rate	Two Year	Five Year	Ten Year
June 5, 2014	0	−10	−1	−4	−1
September 4, 2014	−10	−20	−6	−6	1
December 2, 2015	−20	−30	14	18	20
March 10, 2016	−30	−40	8	10	6

Notes: Changes in yields (in basis points) of German government bonds of different maturities (GBBD02Y, GBBD05Y, GBBD10Y) on days of changes in ECB deposit rate cuts.

by a scalar AR(1) process. But it is straightforward to model it as part of a vector autoregression, as in Wu and Xia (2016). This generalization makes it possible to use the response of the entire term structure to individual announcements to infer parameters of the model. However, it would be necessary to augment the exercise with a description of how $\bar{r}_t$ could change over time and how investors form expectations about future values of $\bar{r}_{t+n}$. This exercise has recently been carried out by Wu and Xia (2017).

3. Empirical Evidence

Table 1 provides some tentative evidence extending the authors' analysis of changes in the European Central Bank's deposit rate into negative territory. Of these four episodes, only the drop in September 2014 looks much like the predicted theory which says that short rates should fall by less than the drop in the policy rate and long rates should fall less than short rates. For the last two episodes, yields on German government bonds actually rose on the days when the ECB cut the deposit rate. Here is the explanation from the *Wall Street Journal* for what happened on March 10, 2016:

> The ECB cut its deposit rate by 0.1 percentage point to minus 0.4% on Thursday, in line with investor expectations. But Mr. Draghi said the ECB had no imminent plans to cut rates further.

Figure 3. Interest Rates in September 2014

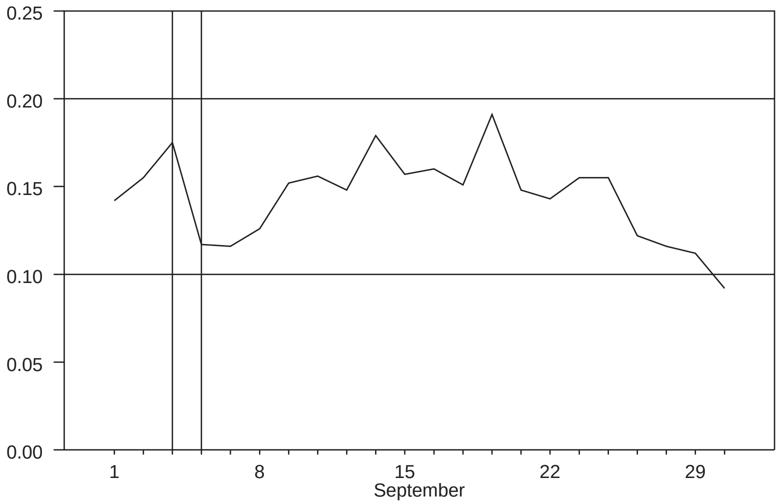

Note: Yield on five-year German bond (in percentage points) after ECB cut policy rate from –0.1 percent to –0.2 percent on September 4, 2014.

> That pushed up short-dated yields, which are particularly sensitive to interest rate moves. Two-year German bond yields, which are particularly sensitive to rate moves, rose sharply to −0.467 percent from −0.553 percent before the announcement.

Just as documented for normal times by Gürkaynak, Sack, and Swanson (2005), the news released to markets by a central bank announcement is more than a one-dimensional object. The market learns not just about the current level of the policy rate but also about its likely future trajectory. This highlights the need for an exercise as in Wu and Xia (2017), in which the shadow rate is part of a vector process and investors form expectations about possible future cuts in the lower bound.

It's also interesting to look in more detail at September 2014, the one episode among these four that seemed most consistent with the theory. As seen in figure 3, the drop in the five-year yield here proved to be temporary, and after a week the yield reached a higher level than it had been before the cut. One possible interpretation is that the cut in the deposit rate lowered the yield by 6 basis points, but subsequent shocks from other sources raised it back again. This is

Figure 4. Federal Reserve's Balance Sheet across Episodes of Quantitative Easing

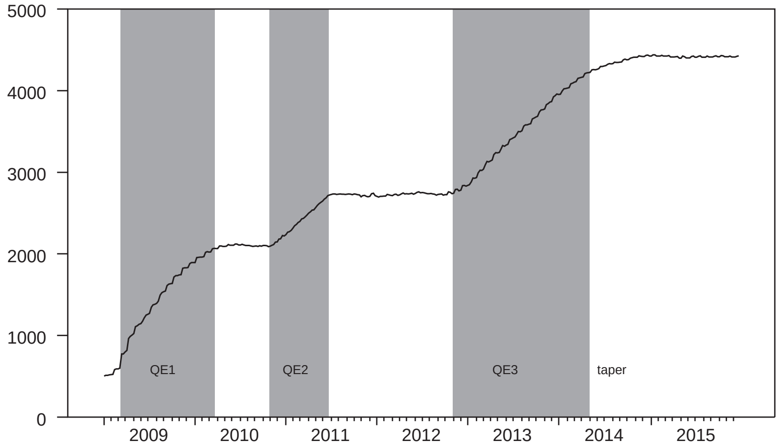

Notes: Federal Reserve holdings of Treasury, agency, and mortgage-backed securities, in billions of dollars, weekly, January 7, 2009 to December 23, 2015. Excludes inflation compensation and unamortized premiums or discounts. Shaded episodes denote March 18, 2009 to March 24, 2010; November 3, 2010 to June 2, 2011; and November 7, 2012 to April 30, 2014.

the implicit interpretation in the many hundreds of papers that have used event-study methodology to evaluate the effects of monetary shocks.

But all of this assumes that the market somehow knows within one day (or within fifteen minutes, for higher-frequency event studies) what the effect will be of a policy instrument that has never been used before. That's not necessarily always going to be the case. We saw a dramatic illustration on November 8, 2016, the day of the U.S. presidential elections. As election returns that night came in with the surprising news that Donald Trump would win the election, there was an immediate dramatic selloff in stock futures, giving an event-study estimate that the election outcome shaved 5 percent off the market value of capital. But the next morning this was all made up and then some. Should the reversal be interpreted as a new shock or as the result of the market continuing to absorb what the election outcome really meant? My view is that nobody knew Tuesday night, or Wednesday morning, and indeed may not know for years, the true implications of a Trump presidency.

Figure 5. Ten-Year Treasury Yields across Episodes of Quantitative Easing

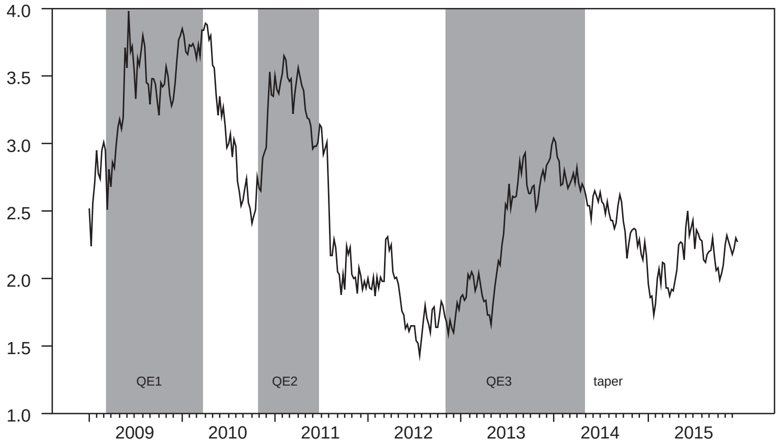

Note: Yields on ten-year U.S. Treasury securities, weekly, January 7, 2009 to December 23, 2005.

This of course is an issue not just with the present paper, but with any studies relying on the event-study methodology. Consider for example the consequences of the three episodes of large-scale asset purchases conducted by the United States, popularly referred to as QE1, QE2, and QE3. These involved major additions to the Federal Reserve's balance sheet, as seen in figure 4. Several announcements were associated with significant drops in the ten-year yield on the day of the announcement, the most dramatic being a 50 basis point decline on March 18, 2009 when the first phase of QE1 was announced. Nevertheless, within a month the yield was back up to the level it had been before the announcement, and at the end of QE1 it stood nearly 100 basis points higher than its value before the program was announced (see figure 5). The event-studies conclusion of Krishnamurthy and Vissing-Jorgensen (2011) that QE1 lowered rates by over 100 basis points is thus implicitly assuming that there were other shocks that in the absence of the program would have raised rates by 200 basis points. Interest rates declined after QE1 ended, started to rise again under QE2, fell after the latter ended, rose under QE3, and fell after tapering of QE3 began, all exactly opposite the conclusions drawn from event studies.

None of this is to question the conventional wisdom that programs like large-scale asset purchases have the potential to lower interest rates in the presence of the zero lower bound, nor Grisse, Krogstrup, and Schumacher's conclusion that negative interest rates are an additional policy tool that could prove effective. Event studies are one of the best ways we have of trying to estimate the effects of these policies, and the current paper makes a useful contribution to the literature. But we should likewise not forget the limitations of this kind of empirical evidence.

References

Christensen, J. H. E., and G. D. Rudebusch. 2015. "Estimating Shadow-Rate Term Structure Models with Near-Zero Yields." *Journal of Financial Econometrics* 13 (2): 226–59.

Gürkaynak, R. S., B. Sack, and E. T. Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking* 1 (1): 55–93.

Hamilton, J. D., and J. C. Wu. 2014. "Risk Premia in Crude Oil Futures Prices." *Journal of International Money and Finance* 42: 9–37.

Krippner, L. 2013. "Measuring the Stance of Monetary Policy in Zero Lower Bound Environments." *Economics Letters* 118 (1): 135–38.

———. 2015. *Zero Lower Bound Term Structure Modeling: A Practitioner's Guide.* Palgrave Macmillan U.S.

Krishnamurthy, A., and A. Vissing-Jorgensen. 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." *Brookings Papers on Economic Activity* (Fall): 215–65.

Ruge-Murcia, F. J. 2006. "The Expectations Hypothesis of the Term Structure When Interest Rates are Close to Zero." *Journal of Monetary Economics* 53 (7): 1409–24.

Wu, J. C., and F. D. Xia. 2016. "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound." *Journal of Money, Credit and Banking* 48 (2–3): 253–91.

———. 2017. "Time-varying Lower Bound of Interest Rates in Europe." Working Paper, University of Chicago.

Capital Injection to Banks versus Debt Relief to Households*

Jinhyuk Yoo
Bank of Korea

I propose a dynamic stochastic general equilibrium model in which the leverage of borrowers as well as banks and housing finance play a crucial role in the model dynamics. The model is used to evaluate the relative effectiveness of a policy to inject capital into banks versus a policy to relieve households of mortgage debt. In normal times, when the economy is near the steady state and policy rates are set according to a Taylor-type rule, capital injections to banks are more effective in stimulating the economy in the long run. However, in the middle of a housing debt crisis, when households are highly leveraged, the short-run output effects of the debt relief are more substantial. When the zero lower bound (ZLB) is additionally considered, the debt relief policy can be much more powerful in boosting the economy both in the short run and in the long run. Moreover, the output effects of the debt relief become increasingly larger, the longer the ZLB is binding.

JEL Codes: E17, E44, E52, E62, G21, H12.

*I thank Byoung Ho Bae, Robert Beyer, Bettina Brüggemann, Cheong-Seok Chang, Albert Marcet, Baptiste Massenot, Atif Mian, Rafael Repullo, Ctirad Slavík, Volker Wieland, and John Williams for instrumental comments and suggestions. I also thank participants of seminars at the Goethe University Frankfurt, the MACFINROBODS Second Consortium Scientific Workshop 2015, Research Department at the Bank of Korea, the 2016 Asia Meeting of the Econometric Society, and the 2016 Conference of the International Journal of Central Banking, "Challenges to Financial Stability in a Low Interest World," for useful feedback. I gratefully acknowledge support from the IMFS and the FP7 Research and Innovation Funding Program (grant FP7-SSH-2013-2). The views expressed in this paper are those of the author and should not be attributed to the Bank of Korea. Author contact: jhyoo@bok.or.kr. Mailing address: Namdaemunno 39, Jung-gu, Seoul, 04531, Korea.

1. Introduction

The Great Recession, which was the largest and longest economic downturn in the postwar era of the United States, was triggered and intensified by the housing debt crisis known as the subprime mortgage crisis.[1] House prices adjusted by the GDP deflator reached a peak in the first quarter of 2006. At the same time, mortgage delinquency rates started to rise from a historically low level. House prices plummeted from the second quarter of 2007 onwards. Losses for financial institutions materialized and credit spreads began to rise.[2] Alongside these financial developments, private consumption slowed down and became very weak from the fourth quarter of 2007, the official beginning of the Great Recession. Non-residential investment started to decline from the next quarter. Finally, after the collapse of Lehman Brothers in September 2008, real GDP in the fourth quarter of 2008 fell drastically by 8.2 percent in annual terms.

It may not be surprising that a collapse in house prices can put the economy into a severe recession through the interaction between the real and financial sector. Housing finance has played a prominent role in advanced economies. According to recent empirical work by Jordà, Schularick, and Taylor (2014), bank loans backed by real estate consisted of roughly 60 percent of total bank lending in 2010, compared with around 30 percent in the 1950s.[3] They find that a rapid increase of home mortgages has mainly contributed to this substantial change in the lending business of banking. In addition, most of the banking crises in advanced economies were associated with boom-bust cycles in house prices. Reinhart and Rogoff (2009) show that five major banking crises during the second half of the twentieth century shared a common pattern:[4] a surge of house prices in the

[1]According to the NBER's Business Cycle Dating Committee, the Great Recession started at 2007:Q4 and ended at 2009:Q2. During the period, real GDP fell by 4.2 percent.

[2]In 2007:Q2, New Century Financial Corporation, a leading subprime mortgage lender in the United States, filed for bankruptcy. Concurrently, the charge-off rate on home mortgages for all U.S. commercial banks started to rise rapidly.

[3]Jordà, Schularick, and Taylor (2014) construct a long-run data set on a wide range of private credit that includes credit to households and credit to firms by commercial banks as well as by other financial institutions such as saving banks, credit unions, and building societies.

[4]Spain (1977), Norway (1987), Finland (1991), Sweden (1991), and Japan (1992).

run-up to a crisis is followed by a sharp decline in the crisis year and in subsequent years together with a prolonged deep recession.

As one of the policy measures to mitigate the severity of the housing debt crisis and ensuing deep economic downturn, the U.S. government promptly used about $500 billion, 3.4 percent of 2008 GDP, to support the U.S. financial sector. The government injected capital worth $245 billion into the U.S. banking sector through the Troubled Asset Relief Program (TARP).[5] It rescued American International Group (AIG), one of the world's major insurance companies, with $67.8 billion of the TARP funds through Treasury purchase of AIG preferred equity, aside from Federal Reserve loans to AIG of maximum $116.8 billion.[6] In addition to financial rescues through the TARP, the housing government-sponsored enterprises, Fannie Mae and Freddie Mac, were nationalized one week before the Lehman Brothers' collapse. The U.S. government committed to putting up to $200 billion into each company and actually injected $187.5 billion capital into the two companies to cover their losses.

For households, the U.S. government pledged only $37.5 billion to refinance home mortgages of those who were in a negative equity position due to the sharp decline in house prices. This amount was tiny compared with the funds to support the financial sector. On top of that, less than one-half of the pledged funds, $18 billion, have been actually spent. Table 1 reports this stark contrast between the two types of post-crisis government interventions.

Many prominent economists, such as Geanakoplos (2010), Stiglitz (2010), Shiller (2012), and Mian and Sufi (2014), criticize the approach biased toward the rescue of financial institutions and argue that more grants for household debt reduction would have provided a significant boost to the economy lacking aggregate demand. Mian and Sufi (2014) claim that the biggest policy mistake of the Great Recession was not to push for mortgage write-downs more

---

[5]The Emergency Economic Stabilization Act authorized the U.S. Treasury to purchase "troubled assets" worth $700 billion in October 2008. Part of the TARP funds, $245 billion, were used to increase banks' capital. Most banks received funds through the Capital Purchase Program, while Bank of America and Citigroup additionally received $20 billion each under the Targeted Investment Program.

[6]The Treasury purchased AIG preferred stock twice: the first purchase was $40 billion in November 2008 and the second was $27.8 billion in January 2009.

Table 1. Comparison of Financial-Sector Support Programs and Household Support Programs during the Great Recession

	Amounts ($Billions)	% of 2008 GDP
Financial-Sector Support Programs	500.2	3.4%
TARP: Capital Purchase Program, etc.	244.9	1.66%
TARP: AIG	67.8	0.46%
Fannie Mae and Freddie Mac	187.5	1.28%
TARP: Mortgage Refinance Programs	18.0	0.12%

Source: Monthly TARP Update for October 1, 2015 and the Bailout Scorecard (projects.propublica.org/bailout).

Note: Mortgage refinance programs include the Making Home Affordable Program ($12.2 billion), the HFA Hardest Hit Fund ($5.7 billion), and the FHA Short Refinance ($0.02 billion).

aggressively. In contrast, leading policymakers at that time, such as Geithner (2014), Summers (2014), and Bernanke (2015), defend the decisions to give priority to financial rescues and emphasize the importance of the credit channel working through financial institutions, though they also admit that more active policies regarding household debt and foreclosure could have been beneficial.

In this paper, I propose a dynamic stochastic general equilibrium (DSGE) model in which the leverage of borrowers as well as banks and housing finance play a crucial role in model dynamics. The model is then used to evaluate the relative effectiveness of a policy to inject capital into the financial sector versus a policy to relieve households of mortgage debt.

The model combines several macrofinancial linkages identified in the literature.[7] It includes two types of households, entrepreneurs and banks. Heterogeneity in the household sector is introduced à la Iacoviello (2005): one household type is more patient than the

[7]Wieland et al. (2016) classify the recent development of macrofinancial DSGE models into three categories: financial accelerator models for corporate investment financing, models with housing finance, and models with an explicit role of banks. I combine three different modeling approaches emphasizing a role of housing finance.

other. In equilibrium, patient households become savers and ultimately supply funds to the economy, while impatient households are borrowers. The impatient households and entrepreneurs borrow funds from banks using real estate as collateral. My model deviates from the standard housing finance models in which borrowing is restricted to a certain fraction of collateral and there is no default. I introduce an agency problem between borrowers and banks by using the costly state verification (CSV) setup of Gale and Hellwig (1985). It implies that the model allows for default in equilibrium. Unlike Bernanke, Gertler, and Gilchrist (1999), I assume that the contractual interest rates are predetermined rather than state contingent. Accordingly, banks make zero expected profits in the perfectly competitive retail loan market, but ex post profits mostly differ from zero. In my model, banks face a leverage constraint making the deviation of the leverage ratio from its target costly as in Gerali et al. (2010). With this leverage constraint, realized profits or losses can affect credit supply. The financial frictions described above are embedded into an otherwise standard New Keynesian model with price and nominal wages rigidities.

The relationship between interest rate spreads and the related leverage ratios can describe key macrofinancial linkages in the model. The risky debt contracts imply that the interest rate spread of each contractual loan rate over the wholesale loan rate, the rate that serves as a benchmark in the retail lending business, positively depends on the leverage of each borrower. Similarly, the bank's optimal decision shows that the interest rate spread of the wholesale loan rate over the deposit rate positively varies with the bank's leverage position. For example, when the leverage of impatient households decreases for some reason, the lending rate spread of home mortgages also narrows, reflecting a decline in default. Faced with lower funding costs, borrower households increase consumption. Meanwhile, realized bank profits due to lower default costs help expanding credit availability. All other things being equal, it further boosts the expenditure of credit-constrained agents, impatient households, and entrepreneurs.

Having in mind that most of the funds to support the financial sector presented in table 1 were injected or committed in one or two quarters after the announcement, I model each policy as a one-time transfer from credit-unconstrained (patient) households to either banks or credit-constrained (impatient) households in policy experiments. The capital injection to banks increases the current period's net worth of banks, while the debt relief to credit-constrained households reduces the outstanding home mortgages. The main findings from the policy experiments are the following

When the economy is near the steady state and policy rates are set according to a Taylor-type rule, the capital injection to banks is more effective in stimulating the economy over the long run. Even though the debt relief to credit-constrained households has a stronger effect on output for the first year, the capital injection policy dominates from the second year onward. The capital injections lead investment to increase, which in turn expands production capacity and results in lower inflation. On the contrary, the debt relief is inflationary and calls for an increase of the policy rate, which reduces investment and the consumption of credit-unconstrained households. In the middle of the housing debt crisis, however, a debt relief policy can be more effective. This is because in such a highly leveraged situation, this policy can reduce the default risk posed by high leverage to a greater extent, thereby resulting in a lower lending rate spread of home mortgages, smaller wasteful foreclosure costs, and a greater short-run stimulus for consumption.

When in addition the zero lower bound (ZLB) constraint is considered, both policies give rise to larger effects on output and help the economy to escape from a liquidity trap earlier than it would without any policy.[8] More interestingly, the effects of the debt relief policy are magnified. The policy-induced inflation under the ZLB constraint leads to a lower real interest rate. The decrease in the real rate boosts investment as well as consumption, or at least significantly weakens crowding-out effects. Therefore, in this environment the debt relief can be much more effective in stimulating the economy both in the short run and in the long run. Moreover, the effects of the debt relief policy on output become increasingly larger as the number of periods that the policy rate is constrained at zero increases.

My model builds on a large literature incorporating financial frictions into a DSGE model, including the prominent groundwork such as Carlstrom and Fuerst (1997), Kiyotaki and Moore (1997),

[8] A liquidity trap is usually defined as the situation in which policy rates cannot fall below zero given that hoarding cash offers an alternative to holding deposit.

and Bernanke, Gertler, and Gilchrist (1999). This earlier work and most of the subsequent research focus on an agency problem between financial intermediaries and their borrowers. These kinds of financial frictions imply that the balance sheets of borrowers become a key factor to explain macrofinancial linkages by affecting credit demand. Recently, in particular after the recent global financial crisis, there has been a growing literature focusing on the agency problem between financial intermediaries and their creditors (e.g., Gertler and Karadi 2011, Christiano and Ikeda 2013, and Kiley and Sim 2014). In those approaches, the balance sheets of financial institutions play an important role in real economy by shifting credit supply. This paper contributes to another growing literature that considers financial frictions in both credit demand and credit supply at the same time and puts an emphasis on their interactions. It includes Gerali et al. (2010), Benes and Kumhof (2011), Clerc et al. (2015), and Iacoviello (2015).

This paper is also related to recent work analyzing the macroeconomic effects of housing prices during the Great Recession. Liu, Wang, and Zha (2013) and Guerrieri and Iacoviello (2015a) find that a collapse in housing prices can explain most of the sharp decline in aggregate demand, while Justiniano, Primiceri, and Tambalotti (2015) argue that such fall in housing prices was not enough to put the economy into a deep recession. As for the modeling of home mortgage default, I follow recent approaches that model default as a put option and impose no direct foreclosure costs on borrower households (for instance, Forlati and Lambertini 2011, Jeske, Krueger, and Mitman 2013, Landvoigt 2014, and Quint and Rabanal 2014).

Regarding policy experiments, a number of papers analyze the macroeconomic consequences of injecting more capital into banks using a DSGE model with a leverage constraint for banks (see Hirakata, Sudo, and Ueda 2013, Kollmann et al. 2013, Kiley and Sim 2014, van der Kwaak and van Wijnbergen 2014, Guerrieri et al. 2015, etc.) Most of them find that the capital injection policy has positive effects on output since it increases credit supply to the productive but credit-constrained sector. In contrast, only a few investigate the macroeconomic effects of reducing household debt. Guerrieri and Iacoviello (2015a) analyze the effects of a lump-sum transfer from credit-unconstrained households to credit-constrained households using a DSGE model with the presence of an occasionally binding constraint, and find that such a transfer can have sizable effects on output when the borrowing constraint binds. Mian and Sufi (2014) estimate the macroeconomic effects of the introduction of shared-responsibility mortgages, which in essence feature an automatic principal reduction when housing prices decline below the purchasing level.[9] They put into perspective several empirical studies such as Mian, Rao, and Sufi (2013), Nakamura and Steinsson (2014), and Mian, Sufi, and Trebbi (2015). They argue that the output effects of new mortgage contracts would be large enough to substantially reduce the severity of the recession. Dogra (2014) uses a simple model in which the economy hits the ZLB by household deleveraging and analyzes the effects of debt relief modeled as a targeted transfer. He finds that debt relief stimulates the economy, but the anticipation of debt relief leads to overborrowing. He shows, nevertheless, that optimal policy is still involved in the use of debt relief up to a certain level.

My contributions to the literature are, first, to design a rigorous policy experiment to reduce households' debt in a structural macro model and to compare the relative effectiveness of this policy and the policy to increase banking capital. Second, as an additional novelty, this paper evaluates those policies considering the zero lower bound, following a new literature to assess fiscal stimulus with the assumption of monetary accommodation (see Cogan et al. 2010, Coenen et al. 2012, and Eggertsson and Krugman 2012). Lastly but not least, the model proposed in this paper applies the financial accelerator mechanism to business loan contracts collateralized by commercial real estate. With this modeling device the model implies that a decline in housing prices leads to a decrease in non-residential investment as well as a rise in default on business loans, which is consistent with the empirical evidence. If a standard modeling setup for housing as in Iacoviello and Neri (2010) is simply combined with a standard risky debt contract for corporate financing as in Bernanke, Gertler, and Gilchrist (1999), then the model implies that a fall in housing prices accompanies a business investment boom.

[9]Shared-responsibility mortgages have two important features that are distinct from the existing ones: the bank offers the protection to borrowers when house prices decline below the purchasing price, while the bank obtains 5 percent capital gains when house prices increase over the purchasing level.

The remainder of this paper is organized as follows. Section 2 describes the model economy and defines the competitive equilibrium. Section 3 explains the calibration strategy and its results. Section 4 analyzes the results from a series of simulations, and section 5 concludes.

2. The Model Economy

Time is discrete and quarterly. The economy is populated by a continuum of two types of infinitely lived households. Each household has unit mass. They differ in their discount factors. One type is more patient than the other. A household obtains utility from consumption and housing services and disutility from labor. Within the household, perfect risk sharing is provided to its members. The nominal wages are set by each type's monopolistic labor union. Each of the patient households has a large number of entrepreneurs, and owns banks, intermediate goods producing firms, retail firms, and capital goods producing firms.

Banks channel funds from patient households (and the banks' own net worth) to impatient households and entrepreneurs. Each bank consists of a wholesale branch and two retail branches. One retail branch deals with home mortgages, while the other handles business loans. The wholesale branch issues wholesale loans to the two retail branches subject to a leverage constraint such that it pays a pecuniary cost for the deviation of the bank net-worth-to-asset ratio from its target. When it comes to the retail lending business, an agency problem arises because the return to the underlying assets posted as collateral is subject to idiosyncratic risk and the realization of an individual shock can only be observed by the bank after paying some cost. Consequently, each retail branch makes an ex ante risky debt contract with its borrowers. Bank profits or losses are accumulated into its net worth after a fraction of the net worth is transferred to patient households.

Entrepreneurs combine loans with their net worth to purchase raw non-residential capital and new residential capital. Then they provide the composite capital services to intermediate goods producing firms that use them with two types of labor to produce intermediate goods. Retail firms operate under monopolistic competition and are subject to implicit costs to adjust nominal prices following

Figure 1. Model Structure

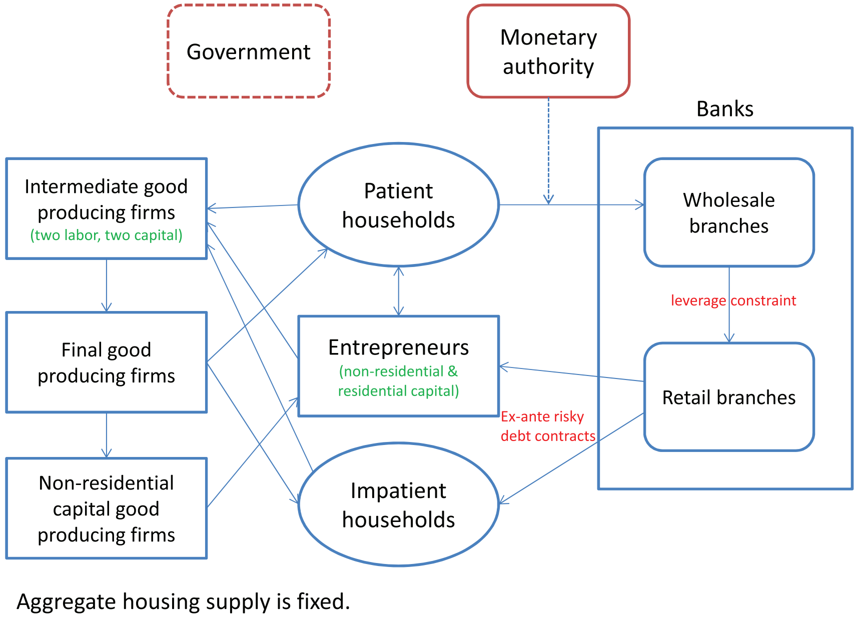

Calvo-style contracts. The constant-elasticity-of-substitution (CES) aggregates of these goods are converted into homogenous final goods. Capital goods producing firms purchase previously installed depreciated capital from entrepreneurs and investment goods from final good producing firms, and produce new installed capital subject to investment adjustment costs.

Aggregate housing supply is assumed to be fixed.[10] The central bank sets the nominal risk-free interest rate according to a Taylor-type rule. The government can collect lump-sum taxes from patient households and give them out to other agents. The structure of the model is depicted in figure 1. In the following, I describe the decision problems of each agent and define the competitive equilibrium of the model economy.

[10]Basically, I exclude residential investment and its spillover effects on the broad economy from the analysis.

2.1 Households

The household sector is composed of two types of households. The discount factor of patient households is higher than that of impatient households ($\beta^P > \beta^I$). Each household ($s \in \{P, I\}$) maximizes the expected discounted sum of per-period utility:

$$V_H^s = E_0 \sum_{t=0}^{\infty} (\beta^s)^t \left\{ \Gamma^s \log(C_t^s - \epsilon C_{t-1}^s) + \chi_t^s \log H_t^s - \psi_l \frac{(L_t^s)^{1+\varphi}}{1+\varphi} \right\}. \tag{1}$$

C_t^s denotes consumption, H_t^s refers to the housing stock owned by each household, and L_t^s denotes labor supplied. I assume habit formation in consumption. As in Iacoviello (2005), utility from housing services is proportionate to the housing stock, and utility is separable in consumption and housing. Γ^s is used for normalization such that the marginal utilities of consumption at the non-stochastic steady state are the inverse of consumption: $\Gamma^s = \frac{1-\epsilon}{1-\beta^s \epsilon}$. I allow for the possibility that each type of household values one unit of housing services differently. Housing preferences are subject to shocks. A decrease in χ_t^s moves preferences away from housing and towards consumption and leisure, so that housing demand decreases and, in the end, housing prices fall.[11] The shock process for each type is given by

$$\log(\chi_t^s) = (1-\rho_\chi)\log(\chi^s) + \rho_\chi \log(\chi_{t-1}^s) + \epsilon_{\chi,t} \quad s \in \{P, I\}. \tag{2}$$

In equilibrium, patient households are savers and impatient households are borrowers. For simplicity, I describe each type's decision problem, taking these equilibrium outcomes into account.

2.1.1 Patient Households

The budget constraint of patient households in real terms is

$$C_t^P + Q_{h,t}(H_t^P - H_{t-1}^P) + D_t = \frac{w_t^P L_t^P}{x_{w,t}^P} + R_{t-1}D_{t-1} + \Upsilon_t^P. \tag{3}$$

[11] For this reason, the shock on housing preferences is also called a housing demand shock.

$Q_{h,t}$ stands for housing prices, D_t is a one-period deposit at period t, R_{t-1} is the real gross interest rate on the last period's deposit D_{t-1}, w_t^P is the real wage rate for patient workers, and $x_{w,t}^P$ is the wage markup between the wage paid by intermediate goods producing firms and the wage paid to households, which accrues to the respective labor union. Υ_t^P stands for profits and net transfers. It includes profits from intermediate goods producers and the labor union, net transfers from banks and entrepreneurs, and lump-sum taxes to government. Patient households earn labor and non-labor income and spend it to consume, to purchase housing, and to save. They maximize their lifetime utility, equation (1), by choosing $\{C_t^P, H_t^P, D_t, L_t^P\}$ given the budget constraint, equation (3).

2.1.2 Impatient Households

The budget constraint of impatient households in real terms is given by

$$\begin{aligned} &C_t^I + Q_{h,t}H_t^I + ac_{zh,t} - Z_t^I \\ &\quad = \frac{w_t^I L_t^I}{x_{w,t}^I} + \int_0^\infty \max\{\omega_t Q_{h,t} H_{t-1}^I - R_{t-1}^L Z_{t-1}^I, 0\} dF(\omega_t) + \Upsilon_t^I. \end{aligned} \tag{4}$$

Z_t^I denotes a one-period mortgage loan at period t, and R_{t-1}^L is the real gross lending interest rate on last period's loan Z_{t-1}^I. w_t^I is the real wage rate for impatient workers, and $x_{w,t}^I$ is the wage markup that the labor union of impatient workers charges. Υ_t^I is profits from the labor union and net lump-sum transfer from government. The term $ac_{zh,t}$ refers to mortgage adjustment costs, and its functional form is quadratic: $ac_{zh,t} = \frac{\kappa_{zh}}{2}\frac{(Z_t^I - Z_{t-1}^I)^2}{Z^I}$. This reflects the fact that home mortgages are, in reality, highly long term and involve many pecuniary and institutional impediments to quick loan adjustments.

I introduce the default risk in housing markets. The value of an individual house is subject to a unit-mean idiosyncratic shock ω_t, which is drawn from the log-normal distribution: $\log(\omega_t) \sim \mathcal{N}(-\frac{\sigma_\omega^2}{2}, \sigma_\omega^2)$. The bank is assumed to offer non-recourse home mortgage loans to the individual members of a household, who use their individual housing as collateral. I further assume that there is no direct cost for households who default. Default, namely, is modeled as a put option. Each member indexed by j decides whether or not to default based on the realization of an individual shock, with the aim to maximize the individual net worth.

$$\max\{\underbrace{\omega_t Q_{h,t} H^I_{t-1}(j) - R^L_{t-1} Z^I_{t-1}(j)}_{\text{No default}}, \underbrace{0}_{\text{Default}}\}$$

The optimal decision rule puts the default threshold $\bar{\omega}_t$ at

$$\bar{\omega}_t = \frac{R^L_{t-1} Z^I_{t-1}(j)}{Q_{h,t} H^I_{t-1}(j)} = \frac{R^L_{t-1} Z^I_{t-1}}{Q_{h,t} H^I_{t-1}} = \frac{m_{t-1}}{\Delta_{h,t}}. \tag{5}$$

The individual members will default on mortgages if mortgage repayment obligations are greater than their housing values. As each member's holdings of mortgages are proportional to that of housing, the index j can be dropped. $\bar{\omega}_t$ can be expressed in terms of the loan-to-value ratio of the previous period, $m_{t-1} = \frac{R^L_{t-1} Z^I_{t-1}}{Q_{h,t-1} H^I_{t-1}}$, and the ex post average return on housing, $\Delta_{h,t} = \frac{Q_{h,t}}{Q_{h,t-1}}$. The default threshold $\bar{\omega}_t$ increases in the household's leverage and decreases in the realized return on housing. Using the default threshold $\bar{\omega}_t$, the budget constraint can be rewritten as

$$C^I_t + Q_{h,t} H^I_t + ac_{zh,t} - Z^I_t \le \frac{w^I_t L^I_t}{x^I_{w,t}} + [1 - \Gamma(\bar{\omega}_t)] Q_{h,t} H^I_{t-1} + \Upsilon^I_t, \tag{6}$$

where $\Gamma(\bar{\omega}_t)$ is the share of the housing value going to the bank.[12] Due to an agency problem, the bank retail branch of home mortgages incurs a cost proportional to the value of foreclosed houses when it forecloses on mortgages. Accordingly, such a risky debt contract must satisfy the following ex ante participation constraint of the bank:

$$E_t\{(\Gamma(\bar{\omega}_{t+1}) - \mu^B G(\bar{\omega}_{t+1})) Q_{h,t+1} H^I_t\} \ge R^r_t Z^I_t. \tag{7}$$

[12] $F(\bar{\omega}_t) = \int_0^{\bar{\omega}_t} dF(\omega^j_t; \sigma_\omega)$ is the foreclosure rate; $G(\bar{\omega}_t) = \int_0^{\bar{\omega}_t} \omega^j_t dF(\omega^j_t; \sigma_\omega)$ denotes a fraction of the foreclosed houses of impatient households; $\Gamma(\bar{\omega}_t)$ are expressed by $[1 - F(\bar{\omega}_t)]\bar{\omega}_t + G(\bar{\omega}_t)$.

μ^B is a foreclosure cost parameter, and R_t^r is the wholesale real lending interest rate that serves as a benchmark rate in retail lending.[13] This constraint states that the expected gain of the bank's contribution to housing investment net of foreclosure cost is at least as high as its funding cost. Finally, impatient households maximize their lifetime utility, equation (1), with respect to $\{C_t^I, H_t^I, Z_t^I, m_t, L_t^I\}$ subject to the budget constraint, equation (6), and the bank's participation constraint, equation (7).

It is worth noting that the participation constraint of the bank holds with equality. It means that the retail branch would make unexpected profits or losses in equilibrium depending on the realization of certain aggregate shocks. Such unexpected profits, $\epsilon_{hb,t}$, are a function of endogenous variables, including the default threshold $\bar{\omega}_t$.

$$\epsilon_{hb,t} = (\Gamma(\bar{\omega}_t) - \mu^B G(\bar{\omega}_t))Q_{h,t}H_{t-1}^I - R_{t-1}^r Z_{t-1}^I \tag{8}$$

In addition, the binding participation constraint of the bank implies that the interest rate spread of the contractual lending rate on mortgages R_t^L over the wholesale lending rate R_t^r is a function of the default threshold $\bar{\omega}_{t+1}$ in expectation.[14] As the right-hand side of equation (9) is an increasing function of $\bar{\omega}_{t+1}$, the interest rate spread rises when the leverage of borrowers increases.[15]

$$\frac{R_t^L}{R_t^r} = E_t \left\{ \frac{\bar{\omega}_{t+1}}{\Gamma(\bar{\omega}_{t+1}) - \mu^B G(\bar{\omega}_{t+1})} \right\} \tag{9}$$

2.1.3 Nominal Wage Decisions

Nominal wage stickiness is introduced in a way analogous to nominal price stickiness as in Smets and Wouters (2007) and Iacoviello and Neri (2010). Each type of household supplies its homogenous labor services to the labor union that serves the interest of each

[13] Alternately, the wholesale lending rate can be thought of as the rate which banks would charge to notional zero-risk borrowers (see Benes and Kumhof 2011).

[14] Plugging equation (5) into equation (7) holding with the equality, we obtain equation (9).

[15] Suppose that $\Omega(x) = \frac{x}{\Gamma(x)-\mu^B G(x)}$. Then $\Omega'(x) = \frac{(1-\mu^B)G(x)+x^2 f(x)}{(\Gamma(x)-\mu^B G(x)6)^2} > 0$ for all $x > 0$. Here, $f(x)$ is a density function of the log-normal distribution.

type.[16] Each union differentiates labor services, sets nominal wages subject to Calvo-style adjustment frictions, and offers labor services to the respective labor packer. Each representative and competitive labor packer aggregates the differentiated labor services into the homogeneous labor services, which are hired by intermediate goods producing firms. The optimal wage rates set by each labor union together with the evolution formula for real wages imply the following wage Phillips curves:

$$\log\left(\frac{\Pi_{w^P,t}}{\Pi}\right) = \beta^P E_t \log\left(\frac{\Pi_{w^P,t+1}}{\Pi}\right) - \kappa_w^P \log\left(\frac{x_{w,t}^P}{x_w^P}\right) \quad (10)$$

$$\log\left(\frac{\Pi_{w^I,t}}{\Pi}\right) = \beta^I E_t \log\left(\frac{\Pi_{w^I,t+1}}{\Pi}\right) - \kappa_w^I \log\left(\frac{x_{w,t}^I}{x_w^I}\right), \quad (11)$$

where $\Pi_{w^P,t} = w_t^P \Pi_t / w_{t-1}^P$ and $\Pi_{w^I,t} = w_t^I \Pi_t / w_{t-1}^I$ refer to wage inflation for patient and impatient households, respectively. $\kappa_w^P = (1-\theta_{w^P})(1-\beta^P\theta_{w^P})/\theta_{w^P}$ and $\kappa_w^I = (1-\theta_{w^I})(1-\beta^I\theta_{w^I})/\theta_{w^I}$ define the slope of each wage equation. Π denotes the non-stochastic steady states of inflation. x_w^P and x_w^I are the wage markup of the patient and impatient households, each.

2.2 Entrepreneurs

Entrepreneurs are modeled in the same way as in Bernanke, Gertler, and Gilchrist (1999) and Christiano, Motto, and Rostagno (2014) with two exceptions. First, the contractual interest rates are predetermined rather than being state contingent. Second, entrepreneurs deal with two types of capital—residential capital and nonresidential capital.[17] Each patient household has a large number of entrepreneurs indexed by j, whose state is summarized by their net worth, $N_t^E(j)$. Each entrepreneur j obtains a loan $Z_t^E(j)$ from the bank's retail branch for business loans and combines it with his net

[16]The essence of wage staggering is to give workers bargaining power to decide wages for a certain period. Equivalently, we can assume that each monopolistic competitive household supplies differentiated labor services to the labor packer and sets nominal wages in a Calvo-style staggering contract.

[17]In this paper, residential capital is used interchangeably with commercial real estate.

worth $N_t^E(j)$ to purchase raw non-residential capital $K_t(j)$ at a price of $Q_{k,t}$ and new residential capital $H_t^E(j) - H_{t-1}^E(j)$ at a price of $Q_{h,t}$. The balance sheet of each entrepreneur at the end of time t is given by

$$Q_{k,t}K_t(j) + Q_{h,t}H_t^E(j) = Z_t^E(j) + N_t^E(j). \tag{12}$$

At period $t+1$, entrepreneurs provide composite capital services, $\tilde{K}_t = H_t^E(j)^{\nu_k} K_t(j)^{1-\nu_k}$, to intermediate goods producers. The return to the composite capital, $\omega_{e,t+1}(R_{t+1}^k Q_{k,t}K_t(j) + R_{t+1}^h Q_{h,t}H_t^E(j))$, is assumed to be sensitive to both idiosyncratic and aggregate shocks. An idiosyncratic shock $\omega_{e,t}$ is modeled to follow a log-normal distribution: $\log \omega_{e,t} \sim N(-\frac{\sigma_{e\omega}^2}{2}, \sigma_{e\omega}^2)$, with $E_t\omega_{e,t} = 1$. The rates of return to non-residential and residential capital are given by

$$R_{t+1}^k = \frac{r_{k,t+1} + (1-\delta_k)Q_{k,t+1}}{Q_{k,t}}, \tag{13}$$

$$R_{t+1}^h = \frac{r_{h,t+1} + Q_{h,t+1}}{Q_{h,t}}. \tag{14}$$

$r_{k,t+1}$ and $r_{h,t+1}$ are competitive market rental rates for non-residential and residential capital, respectively. Non-residential capital depreciates at a quarterly rate of δ_k. As can be seen in equations (13) and (14), an aggregate shock can affect the return to composite capital via either the rental rates or capital prices or both. Note that R_{t+1}^k and R_{t+1}^k are equal across entrepreneurs indexed by j. Similarly to the default threshold of impatient households, equation (5), the decision rule for the entrepreneurs' default threshold is expressed by

$$\bar{\omega}_{e,t+1} = \frac{R_t^E Z_t^E(j)}{R_{t+1}^k Q_{k,t}K_t(j) + R_{t+1}^h Q_{h,t}H_t^E(j)}. \tag{15}$$

R_t^E denotes the real gross lending rate on the business loan Z_t^E. It is worth noting that $\bar{\omega}_{e,t+1}$ is independent of the entrepreneur's net worth and thus her net worth only matters for the size of a certain project. In what follows, I drop the index, j, for notational conveniences. If the realized idiosyncratic shock is below the default threshold, $\omega_{e,t+1} < \bar{\omega}_{e,t+1}$, then entrepreneurs default and hand over all remaining resources to the bank. Meanwhile, the business loan branch has to pay an auditing cost μ^E proportional to the assets of the bankrupt entrepreneurs due to asymmetric information. If $\omega_{e,t+1} \geq \bar{\omega}_{e,t+1}$, entrepreneurs repay debt $R_t^E Z_t^E$ while keeping the surplus return to investment. Therefore, the contractual terms of such a risky debt must satisfy the following ex ante participation constraint of the bank:

$$E_t\Bigg\{(1 - F(\bar{\omega}_{e,t+1}))R_t^E Z_t^E + (1 - \mu^E)\int_0^{\bar{\omega}_{e,t+1}} \omega_{e,t}(R_{t+1}^k Q_{k,t} K_t + R_{t+1}^h Q_{h,t} H_t^E) f(\omega_{e,t}) d\omega_{e,t}\Bigg\} \geq R_t^r Z_t^E. \tag{16}$$

Here, $F(\bar{\omega}_{e,t+1})$ is the default rate for business loans.[18] Finally, each entrepreneur maximizes his share of the expected gross return of capital investment in period $t+1$, $E_t\{(1 - \Gamma(\bar{\omega}_{e,t+1})(R_{t+1}^k Q_{k,t} K_t + R_{t+1}^h Q_{h,t} H_t^E)\}$, subject to equations (12) and (16). I define two variables for entrepreneurial leverage: $\phi_t^k = \frac{Q_{k,t}K_t}{N_t^E}$, $\phi_t^h = \frac{Q_{h,t}H_t^E}{N_t^E}$. Both leverage ratios increase when the entrepreneur's net worth decreases. Each leverage variable also increases when the respective asset prices decline. After some algebra, the entrepreneur's maximization problem can be reformulated as

$$\max_{\phi_t^k, \phi_t^h, \bar{\omega}_{e,t+1}} \quad E_t\{(1 - \Gamma(\bar{\omega}_{e,t+1}))(R_{t+1}^k \phi_t^k + R_{t+1}^h \phi_t^h)\} N_t^E \tag{17}$$

subject to

$$E_t\{[\Gamma(\bar{\omega}_{e,t+1}) - \mu^E G(\bar{\omega}_{e,t+1})](R_{t+1}^k \phi_t^k + R_{t+1}^h \phi_t^h)\} = R_t^r(\phi_t^k + \phi_t^h - 1).$$

The default threshold can be rewritten in terms of leverage variables: $\bar{\omega}_{e,t+1} = \frac{R_t^E(\phi_t^k + \phi_t^h - 1)}{R_{t+1}^k \phi_t^k + R_{t+1}^h \phi_t^h}$. As the first-order conditions imply that $E_t \frac{R_{t+1}^k}{R_t^r} = E_t \frac{R_{t+1}^h}{R_t^r}$, we find that the expectation of $\bar{\omega}_{e,t+1}$

[18]In addition, $\Gamma(\bar{\omega}_{e,t+1})$ denotes a fraction of gross return of capital investment going to the banks, and $G(\bar{\omega}_{e,t+1})$ refers to a fraction of the defaulted value of composite capital. The mathematical expressions are the same as in footnote 12.

increases with total leverage of entrepreneurs, $\phi_t^k + \phi_t^h$.[19] The retail branch for business loans would make unexpected profits in equilibrium according to the realization of a certain aggregate shock. Such profit surprises, $\epsilon_{eb,t}$, are expressed by

$$\begin{aligned}\epsilon_{eb,t} &= (\Gamma(\bar{\omega}_{e,t}) - \mu^E G(\bar{\omega}_{e,t}))(R_t^k Q_{k,t-1} K_{t-1} + R_t^h Q_{h,t-1} H_{t-1}^E) \\ &\quad - R_{t-1}^r Z_{t-1}^E. \end{aligned} \tag{18}$$

Also, reformulations of the participation constraint of banks using the default threshold $\bar{\omega}_{e,t+1}$ show that the lending rate spread of business loans is a function of the default threshold $\bar{\omega}_{e,t+1}$ in expectation. As the right-hand side of equation (19) is an increasing function of $\bar{\omega}_{e,t+1}$, the lending rate spread increases when the entrepreneurial leverage goes up.[20]

$$\frac{R_t^E}{R_t^r} = E_t \left\{ \frac{\bar{\omega}_{e,t+1}}{(\Gamma(\bar{\omega}_{e,t+1}) - \mu^E G(\bar{\omega}_{e,t+1}))} \right\} \tag{19}$$

At the end of the period $t+1$, a fraction $(1-\gamma)$ of each entrepreneur's net worth is transferred to his own household and each entrepreneur receives a lump-sum transfer W^e from the household. Since the ex post net worth of an individual entrepreneur is linear, we can simply integrate individual entrepreneurs' net worth and derive the evolution of total net worth of entrepreneurs as follows:

$$N_t^E = \gamma[(1 - \Gamma(\bar{\omega}_{e,t}))(R_t^k \phi_{t-1}^k + R_t^h \phi_{t-1}^h)\} N_{t-1}^E] + W^e. \tag{20}$$

[19] The first-order conditions are given by

$$\begin{aligned} E_t &\left\{ (1 - \Gamma(\bar{\omega}_{e,t+1})) \frac{R_{t+1}^k}{R_t^r} + \frac{\Gamma'(\bar{\omega}_{e,t+1})}{\Gamma'(\bar{\omega}_{e,t+1}) - \mu^E G'(\bar{\omega}_{e,t+1})} \right. \\ &\quad \left. \times \left[(\Gamma(\bar{\omega}_{e,t+1}) - \mu^E G(\bar{\omega}_{e,t+1})) \frac{R_{t+1}^k}{R_t^r} - 1 \right] \right\} = 0, \\ E_t &\left\{ (1 - \Gamma(\bar{\omega}_{e,t+1})) \frac{R_{t+1}^h}{R_t^r} + \frac{\Gamma'(\bar{\omega}_{e,t+1})}{\Gamma'(\bar{\omega}_{e,t+1}) - \mu^E G'(\bar{\omega}_{e,t+1})} \right. \\ &\quad \left. \times \left[(\Gamma(\bar{\omega}_{e,t+1}) - \mu^E G(\bar{\omega}_{e,t+1})) \frac{R_{t+1}^h}{R_t^r} - 1 \right] \right\} = 0. \end{aligned}$$

With $E_t \frac{R_{t+1}^k}{R_t^r} = E_t \frac{R_{t+1}^h}{R_t^r}$ it holds that $E\{\bar{\omega}_{e,t+1}\} = \frac{R_t^E(\phi_t^k + \phi_t^h - 1)}{R_{t+1}^k \phi_t^k + R_{t+1}^h \phi_t^h} = 1 - \frac{1}{\phi_t^k + \phi_t^h}$.

[20] Refer to footnote 15.

2.3 Banks

Following Gerali et al. (2010), I assume that there is a unit mass of banks and each bank consists of two retail branches and one wholesale branch. The first retail branch is responsible for providing home mortgages Z_t^I to impatient households. The second retail branch gives out business loans Z_t^E to entrepreneurs. I also assume that both retail branches operate in perfect competition. They obtain wholesale loans at the wholesale rate R_t^r and sell them to final borrowers on the condition that the ex ante participation constraint of the bank holds. Therefore, each retail branch sets the contractual lending rate according to equations (9) and (19), respectively. Then the retail branches pass over resulting profits or losses to the wholesale branch.

The wholesale branch has a net worth N_t, which is accumulated out of retained profits, and collects deposits D_t from patient households at the deposit rate R_t. With these funds the wholesale branch issues wholesale loans W_t at the wholesale rate R_t^r while paying a linear operating cost $\kappa_w W_t$ and a quadratic penalty cost when the bank net-worth-to-asset ratio N_t/W_t deviates from its target. The problem of the wholesale branch is to choose loans and deposits to maximize the discounted sum of expected profits subject to the balance sheet constraint:

$$\max_{\{W_t, D_t\}} E_0 \sum_{t=0}^{\infty} (\beta^P)^{t+1} \Lambda_{0,t+1}^P \Pi_{wb,t+1}, \tag{21}$$

where

$$\begin{aligned}\Pi_{wb,t+1} = {} & R_t^r W_t + \Pi_{hb,t+1} + \Pi_{eb,t+1} - R_t D_t \\ & - N_t - \kappa_w W_t - \frac{\phi_n}{2}\left(\frac{N_t}{W_t} - \nu_b\right)^2 N_t\end{aligned}$$

subject to

$$W_t = Z_t^I + Z_t^E = N_t + D_t.$$

Here, $\Pi_{hb,t+1}$ denotes total profits of the home mortgage branch; $\Pi_{eb,t+1}$ refers to those of the business loan branch. Due to the binding participation constraints of the bank, profits or losses resulting from an ex ante risky debt contract, $\epsilon_{hb,t}$ and $\epsilon_{eb,t}$, are an unexpected surprise by construction.[21] To improve empirical validity, I assume that $\epsilon_{hb,t}$ and $\epsilon_{eb,t}$ have persistent effects on the bank's net worth with a decaying factor (ρ_h and ρ_e, respectively), albeit ad hoc.[22] Then $\Pi_{hb,t}$ and $\Pi_{eb,t}$ are expressed by

$$\Pi_{hb,t} = \epsilon_{hb,t} + \rho_h \epsilon_{hb,t-1} + \rho_h^2 \epsilon_{hb,t-2} + \rho_h^3 \epsilon_{hb,t-3} + \cdots ,$$
$$\Pi_{eb,t} = \epsilon_{eb,t} + \rho_e \epsilon_{eb,t-1} + \rho_e^2 \epsilon_{eb,t-2} + \rho_e^3 \epsilon_{eb,t-3} + \cdots .$$

The above equations can be reformulated in autoregressive form:

$$\Pi_{hb,t} = \epsilon_{hb,t} + \rho_h \Pi_{hb,t-1}, \tag{22}$$
$$\Pi_{eb,t} = \epsilon_{eb,t} + \rho_e \Pi_{eb,t-1}. \tag{23}$$

The first-order condition of the wholesale branch's problem is given by

$$R_t^r - R_t = \kappa_w - \phi_n \left(\frac{N_t}{W_t} - \nu_b \right) \left(\frac{N_t}{W_t} \right)^2 . \tag{24}$$

Since the right-hand side of equation (24) is decreasing in the bank net-worth-to-asset ratio around the steady state, the interest rate spread between the wholesale loan rate and the deposit rate increases (up to the first-order approximation) when the bank's leverage goes up. All bank profits are reinvested in banking activity and a fraction (δ_b) of the pre-profit bank net worth is transferred to its own household. Aggregate bank capital evolves according to

$$N_t = (1 - \delta_b) N_{t-1} + \Pi_{wb,t}. \tag{25}$$

[21] $E_t[\epsilon_{hb,t+1}] = 0$, $E_t[\epsilon_{eb,t+1}] = 0$.

[22] I take an approach similar to Guerrieri et al. (2015) that employs an autoregressive process of order 1 for bank losses in their simulations based on an empirically relevant scenario used in the stress tests for the U.S. banking sector. The introduction of the ad hoc adjustment implies a stronger role of the bank's net worth in model dynamics than otherwise. It does not, however, change simulation results qualitatively.

2.4 Production Sector and Nominal Rigidities

In order to introduce price rigidities, I differentiate between competitive intermediate goods producing firms and retail firms. Intermediate goods producing firms hire composite capital services from entrepreneurs and two types of labor from households and solve the following maximization problem given their production technology.

$$\max_{\{L_t^P, L_t^I, K_{t-1}, H_{t-1}^E\}} \left\{ \frac{Y_t}{x_{p,t}} - w_t^P L_t^P - w_t^I L_t^I - r_{k,t} K_{t-1} - r_{h,t} H_{t-1}^E \right\}, \tag{26}$$

where

$$Y_t = [{H_{t-1}^E}^{\nu_k} {K_{t-1}}^{1-\nu_k}]^{\alpha} [{L_t^P}^{\nu_l} {L_t^I}^{(1-\nu_l)}]^{(1-\alpha)}.$$

$x_{p,t}$ is the price markup of final goods over intermediate goods. Retail firms operate in a regime of monopolistic competition and face Calvo-type nominal price frictions. Retailers buy intermediate goods Y_t at the price P_t^w in a competitive market, differentiate the goods, and sell them at price P_t, which includes a markup $x_{p,t} = P_t/P_t^w$ over the marginal cost P_t^w. The CES aggregates of these goods are converted into homogenous final goods, which are purchased by households and capital good producing firms. In every period, each retail firm sets optimal prices with probability $1 - \theta_\pi$ or indexes prices to the steady-state inflation Π with probability θ_π, regardless of the history of its price adjustments. These assumptions deliver the following price Phillips curve:

$$\log\left(\frac{\Pi_t}{\Pi}\right) = \beta^P E_t \log\left(\frac{\Pi_{t+1}}{\Pi}\right) - \kappa_\pi \log\left(\frac{x_{p,t}}{x_p}\right), \tag{27}$$

where $\kappa_\pi = \frac{(1-\theta_\pi)(1-\beta^P\theta_\pi)}{\theta_\pi}$ determines the sensitivity of inflation to changes in the price markup $x_{p,t}$ relative to its steady-state value x_p.

Capital goods producing firms purchase previously installed depreciated capital from entrepreneurs and investment goods I_t from final goods producing firms, and produce new installed capital subject to investment adjustment costs. The capital goods producer solves

$$\max_{I_\tau} E_t \sum_{\tau=t}^{\infty} \Lambda_{t,\tau} \left\{ Q_{k,\tau} I_\tau - \left[1 + s_k\left(\frac{I_\tau}{I_{\tau-1}}\right)\right] I_\tau \right\}, \tag{28}$$

where $s_k(x) = \frac{\kappa_k}{2}(x-1)^2$. The aggregate non-residential capital evolves according to

$$K_{t+1} = (1-\delta_k)K_t + I_t. \tag{29}$$

2.5 Central Bank and Government

The central bank sets the risk-free nominal interest rate based on an interest rate feedback rule that allows for interest rate smoothing and reacts to annual inflation and output.

$$R_t^n = \max\left\{1, (R_{t-1}^n)^{\gamma_R}\left(\frac{\Pi_t^A}{\Pi^A}\right)^{(1-\gamma_R)\gamma_\pi}\left(\frac{Y_t}{Y}\right)^{(1-\gamma_R)\gamma_y}(R^n)^{1-\gamma_R}\right\}, \tag{30}$$

where Π_t^A is year-on-year inflation, which is expressed in quarterly terms as $\Pi_t^A = (P_t/P_{t-4})^{1/4}$. The link between nominal and real interest rates is given by the Fisher equation: $R_t^n = R_t \cdot E_t\Pi_{t+1}$. Government uses fiscal policy only for redistributive purposes: $Tr_t^P + Tr_t^I = 0$. In the baseline model I ignore the redistributive role of the government.

2.6 Competitive Equilibrium

The model is closed by market clearing conditions. As aggregate housing supply is normalized to one, the housing market clears as below:

$$H_t^P + H_t^I + H_t^E = 1. \tag{31}$$

By Walras's law, the good's market clears:

$$Y_t = C_t^P + C_t^I + \left[1 + s_k\left(\frac{I_t}{I_{t-1}}\right)\right]I_t + Adj_t,$$

where $Adj_t = \mu^B G(\bar{\omega}_t;)Q_{h,t}H_{t-1}^I + \mu^E G(\bar{\omega}_{e,t})(R_t^k Q_{k,t-1}K_{t-1} + R_t^h Q_{h,t-1}H_{t-1}^E) + \kappa_w W_t + \frac{\phi_n}{2}(\frac{N_t}{W_t} - \nu_b)^2 N_t + ac_{zh,t}$.

The *competitive equilibrium* consists of a set of prices $\{R_t, R_t^L, R_t^E$, R_t^r, R_t^k, R_t^h, $Q_{h,t}$, $Q_{k,t}$, $r_{h,t}$, $r_{k,t}$, w_t^P, $w_t^I, \Pi_t, x_{p,t}, x_{w,t}^P, x_{w,t}^I\}_{t=0}^\infty$ and a set of real allocations $\{C_t^P, C_t^I, L_t^P, L_t^I, \Lambda_t^P, \Lambda_t^I, H_t^P, H_t^I, H_t^E, D_t, Z_t^I, Z_t^E, K_t, N_t^E, \phi_t^k, \phi_t^h, \bar{\omega}_t, \bar{\omega}_{e,t}, N_t, \Pi_{hb,t}, \Pi_{eb,t}, I_t, Y_t\}_{t=0}^{\infty}$ for a given government policy $\{R_t^n, Tr_t^P, Tr_t^I\}_{t=0}^{\infty}$, a realization of exogenous variables $\{\epsilon_{\chi,t}\}_{t=0}^{\infty}$ and initial conditions $\{H_{-1}^P, H_{-1}^I, H_{-1}^E, D_{-1}, Z_{-1}^I, Z_{-1}^E, K_{-1}, N_{-1}^E, N_{-1}\}$ such that

- households of both types maximize the lifetime utility given the prices;
- each labor union maximizes its profits given the prices;
- the entrepreneurs' allocations solve the problem (17) given the prices;
- the banks' allocations solve the problem (21) given the prices;
- the intermediate goods firms solve the problem (26) given the prices;
- the retail firms maximize their profits given the prices;
- the capital goods producing firms solve the problem (28) given the prices;
- the government budget constraint holds;
- markets clear.

A set of equations describing the equilibrium of the model is summarized in the appendix of Yoo (2017).

3. Calibration

I divide the model parameters into two sets. For the first set of parameters I choose values mostly from the relevant literature. The second set of parameters is endogenously determined in the model to ensure that the model's steady state is consistent with the empirical features related to macro aggregates as well as debt and real estate owned by households and non-financial businesses for the U.S. economy during 1991–2006.

3.1 Exogenously Chosen Parameters

Exogenously chosen parameters are presented in table 2. The discount factor of patient households β^P is set at 0.995, implying that the steady-state risk-free real interest rate is 2 percent. As for other

Table 2. Exogenously Chosen Parameters

Description	Parameter	Value
Households, Production Sector, and Nominal Rigidities		
Discount Factor, Patient Households	β^P	0.995
Habit Formation in Consumption	ε	0.65
Inverse of the Frisch Elasticity	φ	1.0
Disutility Weight on Labor	ψ_l	1.0
Steady-State Price Markup	x_p	1.2
Steady-State Wage Markup	x_w	1.2
Probability of Keeping Prices Fixed	θ	0.85
Probability of Keeping Wages Fixed	θ_w	0.9
Investment Adjustment Cost	κ_k	2.0
Steady-State Inflation	Π	1.005 (2%)
Wage Share of Impatient Households	$1 - v_l$	0.35
Financial Frictions		
Foreclosure Cost	μ^B	0.17
Monitoring Cost	μ^E	0.215
Fraction of Entrepreneurial Net Worth Transferred to Households	$1 - \gamma$	1–0.982
Target Capital-to-Asset Ratio	v_b	0.1
Wholesale Funds Operating Cost	κ_w	0.015/4
Bank Capital Adjustment Cost	ϕ_n	25
Mortgage Adjustment Cost	κ_{zh}	4.0
Autocorrelation of Retail Profits, Home Mortgages	ρ_h	0.9
Autocorrelation of Retail Profits, Business Loans	ρ_e	0.9
Others		
Policy Smoothing Coefficient	ρ_R	0.7
Policy Reaction Coefficient on Inflation	ρ_π	2.0
Policy Reaction Coefficient on Output	ρ_y	0.5/4
AR(1) Coefficient on Housing Preference Shock	ρ_χ	0.97

conventional parameters, I choose standard values in the New Keynesian literature. These include the consumption habit parameter ϵ, the inverse Frisch elasticity of labor supply φ, the steady-state price markup x_p, the steady-state wage markup x_w, and the investment adjustment cost parameter κ_k. The disutility weight on labor steady ψ_l is normalized to 1. I set parameters governing the price and wage rigidities, θ and θ_w, to 0.85 and 0.9, respectively. I choose values from the mid-upper range found in the literature in order to partially compensate the absence of several real or nominal frictions such as variable capital utilization and partial indexation of prices and wages to past inflation. The steady-state annual inflation is set to 2 percent. I set the steady-state wage share of impatient households to 0.35, which is in the mid-range of the existing literature.[23] The resulting income share of these households in aggregate income is 24 percent.

In choosing the parameters related to financial frictions, I rely on a number of previous studies. As in Chatterjee and Eyigungor (2015), the foreclosure cost parameter μ^B is set to 0.17, which means that banks lose 17 percent of the values of houses in case of foreclosure. The parameter for the monitoring cost of business loans μ^E is taken from Christiano, Motto, and Rostagno (2014), who estimate it together with other structural parameters. The parameter governing the transfer from entrepreneurs to their respective households, $1-\gamma$, is determined to be 1–0.982. This lies in between the 1–0.973 value used by Bernanke, Gertler, and Gilchrist (1999) and the 1–0.985 value used by Christiano, Motto, and Rostagno (2014).

Due to the binding participation constraint of the bank, the model-implied profits and losses are unexpected and serially uncorrelated. This implication conflicts with the data that show persistence. To improve the fit of the model, I introduce AR(1) processes (equations (22) and (23)). I set the autocorrelation coefficients, ρ_h and ρ_e, at 0.9, following Guerrieri et al. (2015) who conduct comparative exercises with regard to an exogenous shock to bank capital losses with five structural macrofinancial models.

The parameter for the bank capital adjustment cost ϕ_n is set to 25, which is roughly twice the estimated value of 11.5 in Gerali et al. (2010). It implies that the elasticity of the bank's interest rate spread with respect to the bank net worth is approximately twice as big as that in Gerali et al. (2010).[24] The targeted capital-to-asset ratio ν_b is set to 10 percent. I consider not only the regulatory minimum capital

[23]The wage share of credit-constrained households used in other studies is as follows: 0.36 in Iacoviello (2005), 0.21 in Iacoviello and Neri (2010), 0.42 in Guerrieri and Iacoviello (2015a), etc. Using the Survey of Consumer Finances (SCF), Kaplan and Violante (2014) estimate that between 17.5 percent and 35 percent of U.S. households are liquidity constrained.

[24]Refer to equation (32) in footnote 37.

requirement of 8 percent but also the additional requirement implicitly enforced by market discipline, which seems to be consistent with U.S. data.[25] The operating cost parameter at the wholesale branch of banks, κ_w, is put to 0.015/4, implying that the steady-state interest rate spread of the wholesale loan rate over the deposit rate is 1.5 percent in annual terms. The parameter governing mortgage adjustment costs is set to 4.[26]

As for the interest rate rule's specification, I use 0.7 for interest rate smoothing, 2.0 for the reaction coefficient on inflation, and 0.5/4 for the reaction coefficient on output. Lastly, I employ a fairly persistent AR(1) process for the housing preference shock as in Iacoviello and Neri (2010) and Liu, Wang, and Zha (2013).

3.2 Endogenously Determined Parameters

Endogenously determined parameters are reported in table 3 together with empirical targets used for calibration. The steady-state default threshold for home mortgages $\bar{\omega}$ and the standard deviation of an idiosyncratic shock in households' housing σ_ω are chosen such that the loan-to-value ratio of impatient households is 85 percent and the annual foreclosure rate for home mortgages is 1.5 percent. The former is in line with the average loan-to-value ratios of first-time homebuyers for the 1990s and early 2000s estimated by Duca, Muellbauer, and Murphy (2011), and the latter is the average foreclosure rate for 1991–2006 reported by the Mortgage Bankers Association. In a similar way, the steady-state default threshold for business loans $\bar{\omega_e}$ and the standard deviation of an individual shock in entrepreneurs' capital services σ_{ω_e} are calibrated to agree with both the average of debt-to-net-worth ratio of non-financial business for 1991–2006, 0.42, and the annual business default rate estimated by Christiano, Motto, and Rostagno (2014), 2.25 percent.[27]

[25] The average of the equity-to-asset ratios for the U.S. commercial banks during 2004–06 is 9.9 percent.

[26] I conduct the sensitivity analysis in the appendix of Yoo (2017) since there is no prior information.

[27] Debt of non-financial business is the sum of the credit market instruments of both non-financial corporate business and non-financial non-corporate business. Net worth is the sum of net worth of non-financial corporate business at market prices and net worth of non-financial non-corporate business. All the data are taken from tables B.102 and B.103 in Flow of Funds Accounts.

Table 3. Endogenously Determined Parameters and Selected Steady-State Ratios

Description	Parameter	Value
A. Endogenously Determined Parameters		
Discount Factor, Impatient Households	β^I	0.9598
Steady-State Housing Parameter, Patient Households	χ^P	0.032
Steady-State Housing Parameter, Impatient Households	χ^I	0.190
Capital Share in Production	α	0.307
Commercial Real Estate Share in Production	$1-\nu_k$	0.050
Depreciation Rate	δ_k	0.0226
Steady-State Default Threshold, Home Mortgages	$\bar{\omega}$	0.850
Standard Deviation, Idiosyncratic Shock on Households' Housing	σ_ω	0.060
Steady-State Default Threshold, Loans to Entrepreneurs	$\bar{\omega}_e$	0.316
Standard Deviation, Idiosyncratic Shock on Entrepreneurs' Capital	σ_{ω_e}	0.419
Bank Net Worth Transferred to Patient Households	δ_b	0.005
Transfer Received by Entrepreneurs	W^e	0.0093
Description	**Model**	**Data**
B. Steady-State Ratios		
Targets:		
Ratio of Home Real Estate to Output, $\frac{Q_h(H^P+H^I)}{Y(=C+I)}$	1.7×4	1.7×4
Share of Home Mortgages in Total Loans, $\frac{Z^I}{Z^I+Z^E}$	0.47	0.45
Ratio of Non-residential Fixed Assets to Output, $\frac{K}{Y}$	6.5	6.5
Ratio of Non-residential Investment to Output, $\frac{I}{Y}$	0.147	0.146
Share of Business Residential Fixed Assets, H^E	0.11	0.11
Loan-to-Value Ratio of Credit-Constrained Households, $\frac{R^L Z^I}{Q_h H^I}$	0.85	0.85
Entrepreneurs' Debt-to-Net-Worth Ratio, $\frac{Z^E}{N^E}$	0.46	0.42
Annual Foreclosure Rate, $400F(\bar{\omega})$	1.5%	1.5%
Annual Business Default Rate, $400F(\bar{\omega}_e)$	2.24%	2.25%
Model-Implied Steady-State Ratios:		
Capital-to-Output Ratio, $\frac{K+Q_h H^E}{Y}$	7.3	7.1
Charge-Off Rates, Households, $400(F(\bar{\omega})Z^I-(1-\mu^B)G(\bar{\omega})Q_h H^I)/Z^I$	0.27	0.15
Charge-Off Rates, Entrepreneurs, $400(F(\bar{\omega}_e)Z^E-(1-\mu^E)G(\bar{\omega}_e)(R^k K+R^h Q_h H^E))/Z^E$	0.68	0.74
Households' Aggregate Loan-to-Value Ratio, $\frac{R^L Z^I}{Q_h(H^P+H^I)}$	0.31	0.37
Interest Rate Spread of Home Mortgages in Annual Terms, R^L-R	1.8%p	1.8%p
Interest Rate Spread of Business Loans in Annual Terms, R^E-R	2.2%p	2.1%p
Nominal Risk-Free Rate in Annual Terms, $R^n=R\cdot\Pi$	4.0%	4.1%
Banks' Interest Rate Spread in Annual Terms, R^r-R	1.5%p	–

The capital share in production α, the commercial real estate share in production $1 - \nu_k$, and the depreciation rate δ_k are jointly calibrated to match three sample averages for 1991–2006: the ratio of non-residential fixed assets to output (6.5), the share of business residential fixed assets in total residential fixed assets (11 percent), and the ratio of non-residential investment to output (14.6 percent).[28] The resulting α, $1-\nu_k$, and δ_k are equal to 0.307, 0.050, and 0.0226, respectively.

The parameters governing the steady-state housing weight in utility for the patient and impatient households, χ^P and χ^I, are jointly determined such that the ratio of home real estate to annual output is 1.7 and the share of home mortgages in total loans, home mortgages plus business loans, is 45 percent.[29] The resulting χ^I is 0.190, which is approximately six times as big as the calibrated χ^P of 0.032. These calibration results imply that the impatient households, as a whole, would exert more influence on housing prices than the case when $\chi^P = \chi^I$ is assumed. Guerrieri and Iacoviello (2015a) find that in their housing finance model with the assumption of $\chi^P = \chi^I$, housing services are primarily priced by patient households. Geanakoplos (2010) argues, however, that houses are priced by the most leveraged households because debt enables them to increase their bidding power. The resulting calibration thus helps to capture part of the claim of Geanakoplos (2010).

Lastly, the discount factor of impatient households β^I, the parameter concerning dividends from banks to patient households δ_b, and the transfer received by entrepreneurs W^e are endogenously

[28] To arrive at these numerical values, I first construct an output series, which is consistent with the model definition—the sum of consumption and non-residential investment—using the chain-aggregation methods outlined in Whelan (2002). Then I compute the ratio of non-residential fixed assets to output and the ratio of non-residential investment to output by dividing the respective variables by the model-consistent output. Business residential fixed assets are defined as the sum of ones owned by corporate business and non-corporate business, such as sole proprietorships and partnerships. I compute its share in all residential fixed assets. I use National Economic Accounts (tables 1.1.3, 1.1.5, and 1.1.6) and Fixed Assets Accounts (tables 1.1, 1.2, 5.1, and 5.2) of the Bureau of Economic Analysis.

[29] The data on home real estate and home mortgages are obtained from table B.100 in Flow of Funds Accounts. A model-implied share of home mortgages in total loans ends up with 47 percent at the steady state.

determined to ensure that the related steady-state equations hold. The resulting β^I is 0.9598.

The bottom panel of table 3 reports some selected model-implied steady-state ratios and their empirical counterparts, if they exist. The corresponding data are sample averages computed over the period of 1991–2006. Overall, the model matches quite well many of the empirical moments that are not targeted. They include financial variables such as charge-off rates for home mortgages and business loans, the interest rate spread of the mortgage lending rate over the risk-free rate, and the interest rate spread of the business loan rate over the risk-free rate.[30] The resulting aggregate loan-to-value ratio of households (0.31) is a bit short of the empirical moment (0.37). It could reflect that part of home mortgages can be also issued to patient households, in reality.

4. Results

This section presents the findings of a series of simulation exercises. I begin with an impulse response analysis with two financial shocks: a housing preference shock and a shock to bank losses. Both shocks were identified as critical during the Great Recession. The simulations reveal the key mechanism of the baseline model. The subsequent analysis deals with policy experiments starting from the steady state, that is, policy experiments in non-crisis times. The government levies lump-sum taxes on patient households and uses them either to increase banks' capital or to reduce impatient households' existing debt. Next, I conduct a crisis experiment

[30] Regarding charge-off rates, the corresponding empirical counterparts are charge-off rates on single-family residential mortgages of all commercial banks and those on business loans of all commercial banks, respectively. Charge-off rates are defined as loans removed from banks' balance sheets and charged against loss reserves, net of recoveries as a percentage of average loans in annual terms. I define the mortgage rate spread as the yield on the thirty-year fixed-rate mortgages minus the average yield on the five-year and ten-year Treasury bonds, following Walentin (2014). He proposes this definition based on the observation that the duration of a thirty-year fixed-rate mortgages is, on average, seven to eight years. The interest rate spread of business loans is defined as the difference in the yield on Moody's Baa-rated seasoned corporate bonds and the ten-year Treasury bonds. All the date are obtained from the Federal Reserve Economic Data (FRED) database. More detailed data sources are presented in footnote 42.

to mimic key features of the U.S. housing debt crisis. To be more specific, I feed a sequence of negative housing preference shocks into the model so that the model can replicate the observed decline in housing prices, and take a look at the isolated effects of such a collapse in housing prices on other macro and financial variables. Then, such an environment is used as a laboratory for comparing the consequences of two policies, the capital injection to banks and the debt relief to households. The magnitude and timing of each policy are chosen to be comparable to those of financial-sector support programs that the U.S. government conducted during the recent crisis (see table 1). Last but not least, I analyze how the consideration of the zero lower bound (ZLB) can affect the relative effectiveness of both policies.

In simulating the model, I first compute the non-stochastic steady state where all of the endogenous variables remain constant and the empirical targets are matched. Then I compute the approximated time path of endogenous variables in response to an exogenous shock (or shocks) by log-linearizing the model's equilibrium conditions around the non-stochastic steady state.

4.1 Baseline Simulations

4.1.1 A Housing Preference Shock

Figure 2 presents dynamic responses to a negative housing preference shock that decreases the real housing price on impact by 1 percent. In order to highlight the role of commercial real estate in the model dynamics, figure 2 also contains simulation outcomes from a model without commercial real estate H_t^E. Entrepreneurs in this alternative model economy manage only non-residential capital as in the standard financial accelerator model in Bernanke, Gertler, and Gilchrist (1999).[31] I first look at simulation results from the baseline model, the solid lines, and then compare them with those from the model without H_t^E, the dashed lines.

As seen in panel D, housing prices fall immediately by 1 percent and then very slowly return to the steady state. Given the outstanding debt, the decline in housing prices increases borrowers' leverage.

[31] Technically, I set ν_k at 0.000001, while ignoring a share of business residential capital as a target.

Figure 2. Impulse Responses to a Negative Housing Preference Shock in the Models with and without Commercial Real Estate

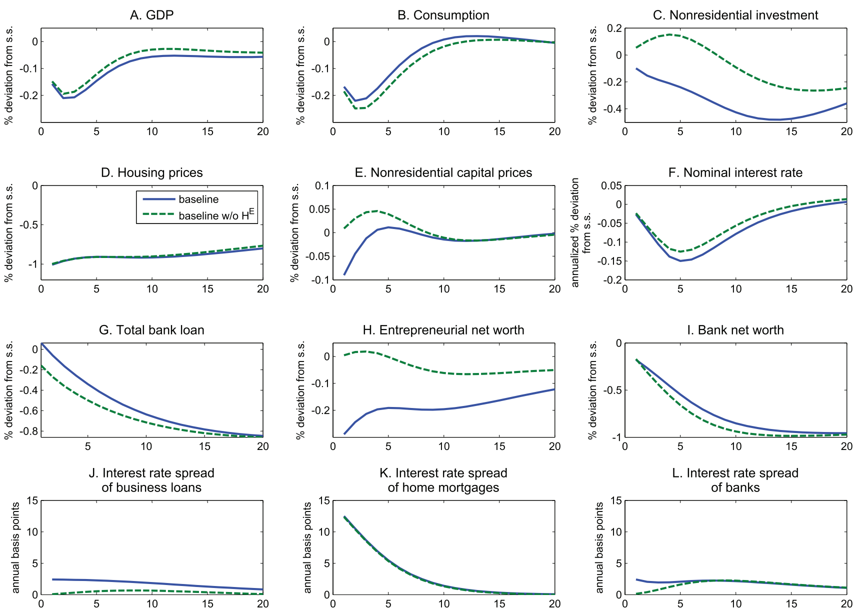

Notes: The horizontal axis represents quarters after the shock. The nominal interest rate is annualized. The simulations show dynamic responses to a housing preference shock that decreases housing prices on impact by 1 percent in both models.

It implies that more members of impatient households and entrepreneurs default on their loans. The increased number of defaults in turn causes losses for banks. In order to satisfy their participation constraint, banks raise the contractual lending rates relative to the wholesale lending rate. The interest rate spread of home mortgages (in panel K) goes up by 12 annual basis points, while that of business loans (in panel L) increases by 2 annual basis points. Since commercial real estate is part of entrepreneurial assets, the effects of housing prices on the entrepreneurial net worth are smaller than those on the net worth of impatient households. Facing higher borrowing costs, borrowers demand less credit. Meanwhile, a fall in housing prices also tightens credit supply, as the bank net worth (in panel I) declines and the bank leverage rises. As a result, the interest rate spread of banks (in panel L) increases by 2 annual basis points.

Given higher borrowing costs, the impatient households significantly cut back on consumption, making aggregate consumption decrease by roughly 0.2 percent (in panel B).[32] The entrepreneurs also reduce demand for non-residential capital, which leads to a decline in investment as well as in its price (see panels C and E). Interestingly, the decrease in housing prices has a persistent effect on non-residential investment. The trough is reached in the third year after the shock. As shown in panel A, GDP, defined as the sum of aggregate consumption and non-residential investment, decreases by about 0.2 percent. In response to the decrease in output and inflation (not shown in figure 2), the central bank decreases its nominal interest rate (in panel F).

Now I compare simulation outcomes from the baseline model with those from a model without residential capital H_t^E. Noticeable differences are found in the dynamic responses of non-residential investment, its prices, and the entrepreneurial net worth (in panels C, E, and H). Since entrepreneurs do not deal with residential capital in the alternative model, housing prices cannot directly affect the entrepreneurs' net worth. Rather, the reduction in the equilibrium real interest rate, which is induced to compensate the reduced consumption of impatient households, boosts investment demand, and thus raises the price of non-residential capital. Non-residential investment and its prices increase for nearly two years before they fall below the steady-state mainly due to the halting credit supply.

The simulation using the baseline model shows that a decline in housing prices leads to a decrease in non-residential investment, which is in line with the empirical evidence presented in Liu, Wang, and Zha (2013). In contrast, the simulation with the alternative model, where a standard modeling setup for housing services like in Iacoviello (2005) or Iacoviello and Neri (2010) is simply combined with a financial accelerator mechanism for new business investment like in Bernanke, Gertler, and Gilchrist (1999), predicts the opposite

[32]The consumption of patient households increases due to a decline in the equilibrium real interest rate.

with respect to housing prices and non-residential investment.[33] Due to a similar logic, a decline in housing prices leads to a business investment boom in a model where credit-constrained households borrow funds against housing collateral and unconstrained households can only access investment technology.[34]

4.1.2 A Shock to Bank Losses

I investigate the effects of capital shortfalls in banks on the economy. Such a disturbance directly exacerbates credit supply conditions. To evaluate the implications of bank capital losses in the baseline model in contrast to those in other macrofinancial models, I conduct the same simulation exercise with a shock to bank losses as in Guerrieri et al. (2015).

They model a shock to bank losses as a lump-sum transfer from the banking sector to households, who are the ultimate suppliers of funds in the economy, and assume that the shock follows an AR(1) process with an autocorrelation coefficient of 0.9. An exogenous disturbance is then fed into the model, so that the banking sector incurs cumulated losses worth 7.5 percent of annual steady-state GDP for nine quarters.[35] Guerrieri et al. (2015) carry out the described simulation with five macroeconomic models, all of which consider an explicit role of banks and have been developed by staff economists at the Federal Reserve Board. They include the model of Iacoviello (2015), the model of Covas and Driscoll (2014), the model of Kiley and Sim (2015), Queralto's model, and the model developed by Guerrieri and Jahan-Parvar. Each model takes a different approach

[33]Kollmann et al. (2013) and Clerc et al. (2015) simply combine the two macrofinancial modeling devices. Regardless of whether or not they recognize their model's implications on housing prices and business investment, they rightly exclude a housing preference shock from a list of exogenous shocks considered.

[34]In the models of Iacoviello and Neri (2010) and Justiniano, Primiceri, and Tambalotti (2015), business investment increases in response to a negative housing preference shock because of these modeling assumptions.

[35]This numerical reference is chosen based on the stress tests for the U.S. banking sector under a severely adverse scenario comparable to the Great Recession, which were conducted for the Comprehensive Capital Analysis and Review (CCAR) of 2013. The amount of the cumulated losses, 7.5 percent of annual GDP, is used to pin down the magnitude of an initial disturbance.

Figure 3. Impulse Responses to a Transfer Shock from the Banking Sector to Patient Households

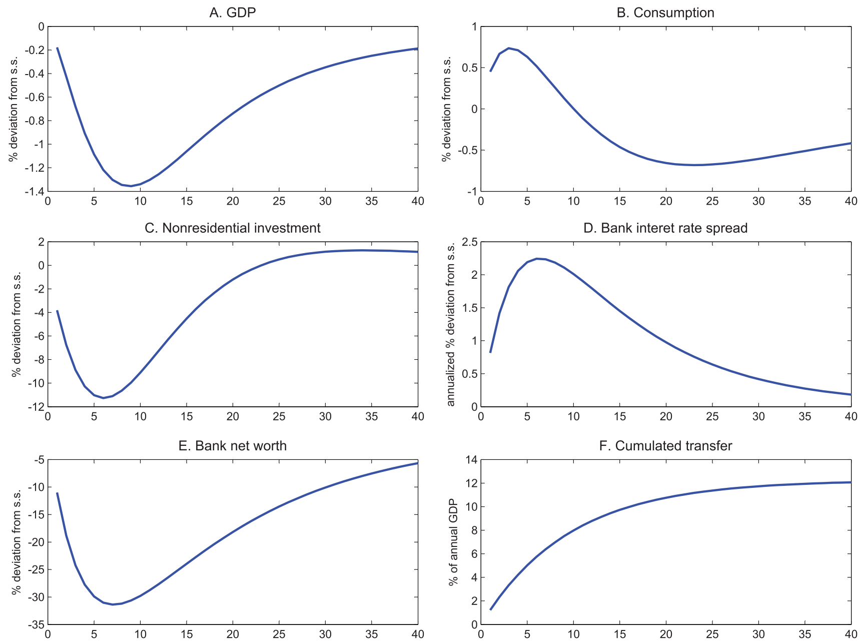

Notes: The horizontal axis represents quarters after the shock. The exogenous shock to transfer resources from the banking sector to patient households is fed into the model, so that the banking sector incurs cumulated losses worth 7.5 percent of annual steady-state GDP for nine quarters. This figure is directly comparable to figure 17 of Guerrieri et al. (2005).

to formulate macrofinancial linkages, but all of them are calibrated or estimated on U.S. data. The authors therefore argue that the model comparison can offer a "model-based confidence interval" with respect to the shock which originated from the banking sector.

Figure 3 presents the implications for a transfer shock from the banking sector to patient households in my model.[36] As the transfer shock causes the bank's net worth to decrease, the bank's interest

[36] To make the direct comparison easier, figure 3 includes the same variables as in figure 17, page 48 of Guerrieri et al. (2015).

rate spread rises for any given level of credit to borrowers.[37] As can be seen in panel E, the bank net worth initially falls by about 10 percent and then declines further by over 30 percent in six quarters due to an endogenous decrease in the bank's retained earnings as well as the persistence of the shock process. Thereafter it gradually returns to the steady state. The bank's interest rate spread shows symmetric dynamics. It increases to more than 2.0 percentage points in annual terms and then steadily returns to the long-run level. With the soaring funding costs, entrepreneurs reduce the demand for capital and, as a result, non-residential investment decreases rapidly, as shown in panel C. For the same reason, impatient households reduce consumption. However, patient households, who receive a persistent wealth transfer from the banking sector and face a lower interest rate on deposits, increase consumption enough to offset the decrease in the consumption of impatient households. The aggregate consumption therefore goes up for the first two years (see panel B). To sum up, the aggregate demand, or GDP, decreases by nearly 1.4 percent in the two years after the initial shock occurs and then gradually returns to the steady-state level.

Compared with the output dynamics of the five macrofinancial models in Guerrieri et al. (2015), GDP in my model lies in the mid-range. This comparison suggests that my model delivers quantitatively acceptable implications for the role of banking capital.

4.1.3 Policy Experiments in Normal Times

So far, I have analyzed dynamic responses to two financial shocks and have found that the leverage of both borrowers and banks plays an important role in the model dynamics. When the leverage of impatient households or entrepreneurs increases, the corresponding lending rate spread also goes up, reflecting a rise in default. Faced with higher funding costs, credit-constrained agents cut back on spending. Meanwhile, when the bank's leverage increases,

[37]The log-linearized first-order condition of the bank around the steady state shows this relationship:

$$\hat{R}_t^r - \hat{R}_t = \frac{\phi_n \nu_b^3}{R}\hat{W}_t - \frac{\phi_n \nu_b^3}{R}\hat{N}_t. \tag{32}$$

the bank reduces loan supply, thereby increasing the wholesale loan rate relative to the rate at which the bank borrows. This adverse credit supply leads credit-dependent agents to refrain from spending.

Against the backdrop of these adverse disturbances, I carry out policy simulations to reduce the leverage of credit-constrained agents. More specifically, I compare the effects of a policy to inject capital into banks with those of a policy to reduce the outstanding debt of impatient households. For simplicity, I model each policy as a one-time transfer worth 1 percent of the steady-state annual GDP from the patient households either to banks or to the impatient households.[38] The transfer to banks simply increases the banks' net worth. The policy to reduce the debt of impatient households needs more explanation. This policy involves two effects on the recipients. It increases their net worth like a simple transfer. On top of that, it scales down their leverage and thus reduces the likelihood of default. Thanks to the reduced foreclosure costs, banks make profits and will charge a lower mortgage rate relative to the wholesale lending rate,

[38] In other words, the government earns funds by levying lump-sum taxes on patient households and uses them to support the financially constrained agents. The mathematical expressions are given as follows. Let $Tr_t = 0.04Y^{GDP}$. For a transfer to banks, Tr_t is simply added to the equation for bank capital accumulations at period t.

$$N_t = (1 - \delta_b)N_{t-1} + \Pi_{wb,t} + Tr_t$$

Regarding debt relief to impatient households, I define an instrument for the debt relief policy, τ_t^r, such that $Tr_t = \tau_{\mathbf{t}}^{\mathbf{r}} R_{t-1}^L Z_{t-1}^I$. The budget constraint of impatient households is rewritten with τ_t^r:

$$\begin{aligned} &C_t^I + Q_{h,t}H_t^I + ac_{zh,t} - Z_t^I \\ &\quad = w_t^I L_t^I + \int_0^\infty \max\{\omega_t Q_{h,t}H_{t-1}^I - (1 - \tau_{\mathbf{t}}^{\mathbf{r}})R_{t-1}^L Z_{t-1}^I, 0\} dF(\omega_t) + \Upsilon_t^I. \end{aligned} \quad (33)$$

The default threshold, $\bar{\omega}_t$ is adjusted by the factor of τ_t^r.

$$\bar{\omega}_t = \frac{(1 - \tau_{\mathbf{t}}^{\mathbf{r}})R_{t-1}^L Z_{t-1}^I}{Q_{h,t}H_{t-1}^I} \quad (34)$$

As long as the policy is unexpected, other optimality conditions of impatient households do not change. A more detailed explanation is provided in the appendix of Yoo (2017).

Figure 4. Effects of One-Time Transfer Policies Worth 1 Percent Steady-State Annual GDP

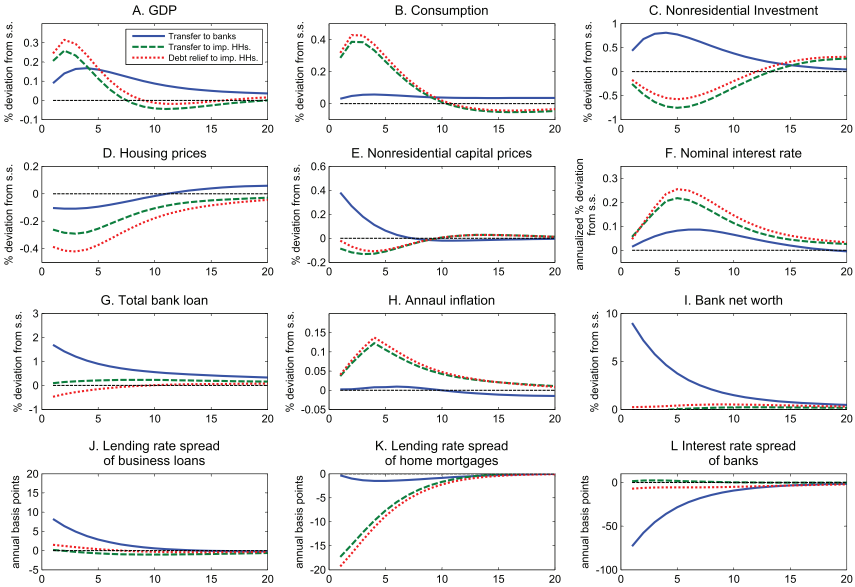

Note: The horizontal axis represents quarters after each policy shock.

which relieves households of an interest repayment burden. Therefore, they have more incentive to increase current consumption. For comparison, I additionally consider a simple lump-sum transfer to impatient households.

Figure 4 shows the dynamic responses to three one-time transfer policies worth 1 percent of steady-state annual GDP starting from the steady state: a transfer to banks (solid line), a debt relief to impatient households (dotted line), and a transfer to impatient households (dashed line). The transfer to banks increases credit availability substantially, so that the interest rate spread of banks decreases by around 70 annual basis points and total bank loans increase on impact by almost 2 percent (see panels L and G). Consequently, consumption and non-residential investment increase. The impact on investment (in panel C) is large. The increased investment demand causes capital prices to rise, and thus entrepreneurial net worth also goes up. This further boosts non-residential investment in the subsequent periods. In contrast, the impact on aggregate consumption is fairly mild. The consumption of impatient households gradually increases partly due to mortgage adjustment costs and is mostly offset in the aggregate by a decrease in the consumption of patient households. Housing prices decline as patient households reduce their demand for housing services due to negative wealth effects. The effects on inflation are negligible and thus the central bank raises the nominal rate only a little.

Overall, real GDP increases by 0.17 percent before slowly returning to its long-term level. The output effects are quite persistent: GDP increases by 0.14 percent for each of the first two years and by 0.08 percent for the third year. The cumulative output increase after ten years amounts to roughly 0.6 percent relative to the baseline annual GDP.

I now turn to evaluating a policy to reduce the debt of impatient households with funds raised from patient households. The marginal propensity to consume of the impatient households is higher than that of the patient households due to a lower discounting factor, so that a simple transfer, by itself, can lead to an increase in aggregate consumption. The dashed line shows the simulation outcomes with the lump-sum transfer to impatient households. Yet, the debt relief has an additional channel to increase consumption. The reduction in existing debt results in a lower funding cost for new home mortgages through the decreased mortgage default risk. The lending rate spread of home mortgages decreases on impact by 20 annual basis points, which is 2 basis points larger relative to those under the policy with a simple transfer. Consequently, aggregate consumption increases initially by 0.3 percent and then goes up by 0.43 percent in the next period, which is 0.05 percent higher than with the policy using a simple transfer.

However, this sharp boost in consumption leads to a rise in inflation. In response to such an inflationary development, the central bank raises the nominal interest rate, which translates into a rise in the real interest rate due to price rigidities. A higher real rate causes patient households to cut back on consumption further, so that ten quarters after the policy is undertaken, aggregate consumption falls below its steady state (see panel B). It also leads entrepreneurs to refrain from investment (in panel C) and thus non-residential capital prices decline (in panel E). It is particularly interesting to see that housing prices (in panel D) decline more under debt relief than under capital injections to banks, even though debt relief increases the impatient households' appetite for housing services. This is because under this policy, patient households reduce housing demand to a greater extent due to higher real interest rates in addition to negative wealth effects.

Overall, real GDP increases sharply by nearly 0.3 percent in the first year and then slows down to 0.1 percent in the second year. Then it falls below the steady state and turns negative in the third year and thereafter returns to the non-policy path. In the case of the policy with a simple transfer, stimulative output effects are more short-lived. The cumulative increase in output after ten years reaches 0.45 percent over baseline annual GDP.

To sum up, the output effects of a debt relief to credit-constrained households are roughly twice as large as those after injecting resources to banks in the first year. However, the latter dominates the former from the second year onward, thereby making a transfer to banks more effective over the long run. It is because such a transfer causes non-residential investment to increase, expanding production capacity and thereby resulting in lower inflation. On the contrary, household debt relief gives rise to higher inflation and therefore leads to the central bank's reaction to raise the interest rate, which significantly dampens investment as well as the consumption of credit-unconstrained households.

4.2 An Application to the U.S. Housing Debt Crisis

4.2.1 The Effects of the Collapse in Housing Prices

In this subsection I analyze the effects of the sharp decline in house prices observed in the U.S. subprime mortgage crisis. House prices started to decline from the first half of 2006, long before the recession began, and plummeted until the second quarter of 2009. Then they remained more or less flat until 2011 before they started to recover. Based on the S&P Case-Shiller U.S. National Home Price Index adjusted by the implicit GDP deflator, house prices dropped by around 30 percent from 2006:Q1 to 2009:Q2. If the stochastic trend is removed, house prices declined by nearly 37 percent relative to the long-run trend over the same periods.[39]

To engineer such a fall in housing prices, I feed a series of negative housing preference shocks into the model, similarly to the approach taken by Justiniano, Primiceri, and Tambalotti (2015). This strategy to generate the decline in housing prices is based on several studies which, using an estimated structural model with housing, document that shocks on housing preferences are the primary determinant of fluctuations in housing prices. Liu, Wang, and Zha (2013) show that housing preference shocks explain around 90 percent of fluctuations in housing prices at all the different forecasting horizons from one quarter to six years.[40] Also, Guerrieri and Iacoviello (2015a) find that the decline in housing prices observed during the Great Recession is almost exclusively explained by the realizations of housing preference shocks.[41] Therefore, this exercise allows us to quantify the isolated effects of the observed decline in housing prices on macroeconomic aggregates as well as the relevant financial variables.

Figure 5 presents the model-simulated outcomes together with their analogous data observed from the first quarter of 2006 to the fourth quarter of 2014. The actual data include GDP, consumption,

---

[39]I use a one-sided Hodrick-Prescott (HP) filter with a value of $\lambda = 100{,}000$ to decompose a time series into a trend component and a cycle component. Guerrieri and Iacoviello (2015a) argue that choosing high λ produces plausible estimates for macroeconomic aggregates and financial variables during the Great Recession. In addition, the one-sided filter prevents the decomposition from being influenced by the correlation between the current observation and the subsequent observations.

[40]The finding is based on variance decompositions reported in table 3 of Liu, Wang, and Zha (2013). Their model includes a representative patient household and a representative impatient entrepreneur. The entrepreneur's borrowing is tied to the value of collateral assets consisting of commercial real estate and non-residential capital. The model employs eight exogenous shock processes: an intertemporal preference shock, a housing preference shock, a labor supply shock, permanent and transitory technology shocks, permanent and transitory investment-specific shocks, and a collateral shock.

[41]The results are obtained from historical decompositions presented in figure 5 of Guerrieri and Iacoviello (2015a). The model includes patient households and impatient households. The debt of impatient households is occasionally tied to their housing value. The model uses six exogenous shock processes: an intertemporal preference shock, a housing preference shock, a price markup shock, a wage markup shock, an investment shock, and a monetary policy shock.

Figure 5. An Application to the U.S. Housing Debt Crisis (2006:Q1–2014:Q4)

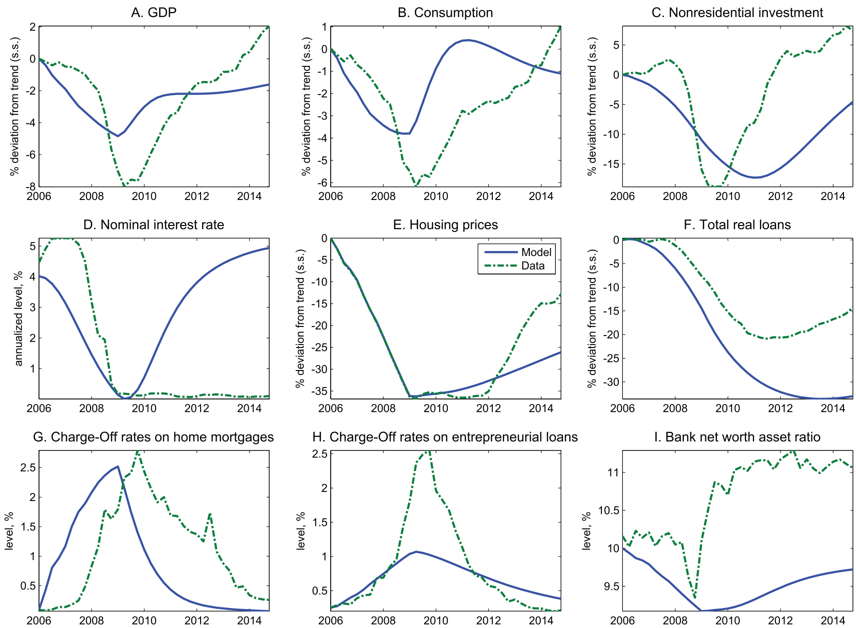

Notes: For the data series, the horizontal axis ranges from 2006:Q1 to 2014:Q4. GDP, consumption, non-residential investment, and total real loans are expressed in real, per capita terms. These quantity variables as well as housing prices are detrended using a one-sided HP filter ($\lambda = 100{,}000$) and are shifted in the vertical axis such that the estimates for 2006:Q1 are at zero. The rest of the data series are shown without transformation. The model-simulated series start from the steady state. Please refer to the text for more details.

non-residential investment, the sum of home mortgages and business credit, housing prices, the federal funds rate, charge-off rates on single-family residential mortgages, charge-off rates on business loans, and total equity to total assets for all commercial banks.[42]

[42]I obtain data series on consumption and non-residential investment from U.S. National Economic Accounts. Using the chain-aggregation methods, I compute the model-consistent output, defined as the sum of consumption and non-residential investment. The rest of the data can be downloaded from FRED (https://research.stlouisfed.org/fred2/): Home Mortgages (HMLBSHNO), Credit Market Instruments for Nonfinanical Corporate Business (TCMILBSNNCB),

The aggregate quantity variables are expressed in real, per capita terms.[43] These variables, in addition to housing prices, are detrended using a one-sided HP filter (with a value of $\lambda = 100{,}000$) in order to remove the low-frequency components.[44] Each data series in figure 5 is normalized such that its estimate for the first quarter of 2006 is at zero. The rest of the variables are displayed without transformation. The model-simulated series starts from their steady-state values. GDP, consumption, non-residential investment, housing prices, and total real loans are expressed in the percentage deviation from their own steady state. The remaining variables are measured in the same way as their data counterpart.[45]

Innovations in housing preferences are introduced into the model for thirteen quarters in a row, such that the model can replicate the historical development of housing prices during the periods from 2006:1Q to 2009:Q2. As can be seen in panel E, the model-implied housing prices do indeed plummet as much as the actual data over the same period. Once no more innovations are fed into the model, the model-implied housing prices gradually recover to the long-run level, which is broadly in line with the actual data. With such a collapse in housing prices, defaults on home mortgages soar and, consequently, so do bank losses. The annualized charge-off rate on home mortgages rises from 0.1 percent to more than 2.5 percent for the initial thirteen quarters (in panel G). This surge accounts for over 90 percent of the peak in the actual charge-off rates. Another observation is that the model-implied charge-off rates are shifted

Credit Market Instruments for Nonfinanical Noncorporate Business (TCMILBSNNB), S&P Case-Shiller U.S. National Home Price Index (CSUSHPISA), Effective Federal Funds Rate (DFF), Charge-Off Rate on Single Family Residential Mortgages (CORSFRMACBS), Charge-Off Rate on Business Loans (CORBLACBS), and Total Equity to Total Assets for Banks (EQTA). In parentheses is an ID for each data series in FRED. Total loans consist of home mortgages and credit market instruments for non-financial corporate business and non-finanical non-corporate business.

[43]The credit variable is deflated by the GDP deflator. The variables are converted into per capita terms by dividing each by the working-age population, which is obtained from the Organisation for Economic Cooperation and Development (OECD).

[44]I use quarterly data covering the period from 1985:Q1 to 2014:Q4.

[45]For a better comparison, the model simulations for the two charge-off rates start from the level of 2006:Q1, not from the steady state.

to the left by around four to six quarters relative to the actual data. This is mainly because in the model, the foreclosure process is completed within a period, but in reality it takes one year or more for financial institutions to foreclose home mortgages. Taking into account an actual time lag for the foreclosure completion, the model-implied charge-off rates on home mortgages explain well the actual outcomes. The differences between the model-implied and empirical consumption mirror those between the model-implied and empirical charge-off rates, because a rise in defaults reduces consumption through the resulting higher lending costs. A comparison of consumption from the model and data suggests that the observed housing prices explain roughly 60 percent of the actual decline in consumption.

The collapse in housing prices also increases defaults on entrepreneurial loans, thereby raising those loans' charge-off rates. Panel H suggests that other shocks must have increased charge-off rates on business loans during the Great Recession. One of them could be a shock to the second moment of an idiosyncratic shock process influencing the return to entrepreneurial capital as suggested in Christiano, Motto, and Rostagno (2014). Meanwhile, the observed decline in housing prices results in a delayed but substantial decrease in nonresidential investment due to the reduction of both credit demand and credit supply. Overall, the model predicts that lower housing prices observed during the periods from 2006:Q1 to 2009:Q2 are associated with a decrease in aggregate output below 5 percent. This accounts for 60 percent of the decline in output from the actual data.

In response to the development of GDP and inflation (not shown in figure 5) induced by the collapse in housing prices, the central bank reduces the nominal interest rate by 4 percentage points. Starting at a steady-state level of 4 percent, the nominal rate almost reaches zero (0.1 percent) thirteen quarters after the initial innovation occurs. As losses in banks are accumulated, banks' net-worth-to-asset ratio ratios decline to roughly 9.2 percent. The drop in the model-implied bank net-worth-to-asset ratio is comparable to that of its empirical counterpart. While the actual bank capital-to-asset ratio went up to over 11 percent from 2008:Q4 after banks received the TARP funds from the government and stayed over 11 percent, the model predicts a fairly gradual returning to the long-run level. These differences occur mainly due to the model's inability to account for the post-crisis government responses and their impact on the value of banks' marketable assets. As for total real loans, we find that the deleveraging implied by the model is much faster than what had happened during the Great Recession.

4.2.2 Policy Experiments in the Middle of the U.S. Housing Debt Crisis

As presented in the introduction, the U.S. government used about $500 billion to recapitalize the U.S. financial system. This amount of funds corresponds to 4.2 percent of the sum of 2008 annual consumption and non-residential investment. In this subsection, I analyze the effects of this recapitalization program and compare them with those of a debt relief program providing the same amount of funds. With the observation that almost all of the funds to support the financial sector were disbursed or committed in one or two quarters after the Troubled Assets Relief Program (TARP) was set up in September 2008, I model a policy intervention as a one-time unexpected shock, namely, a one-shot transfer worth 4.2 percent of annual GDP (consumption + non-residential investment) from patient households to either banks or impatient households.[46] To focus on the effects caused by the policies, I abstract from a distortionary fiscal policy and assume that Ricardian equivalence holds.[47]

The results from the preceding policy exercises in normal times (in section 4.1.3) imply that the bank recapitalization of 4.2 percent of annual GDP would increase GDP by nearly 1.2 percent for the first two years, and by 1.3 percent for the next eight years. In the case of the debt relief to impatient households, GDP would increase

[46]Ninety-nine percent of the TARP funds used to support the financial sector was disbursed or pledged from 2008:Q4 to 2009:Q1. The maximum commitment to rescuing Fannie Mae and Freddie Mac was initially announced in 2008:Q4 and was revised once in 2009:Q1.

[47]All lump-sum taxes are assumed to be paid by patient households. As they behave according to the permanent income hypothesis, the timing of taxes is irrelevant for their consumption-saving decisions.

by 1.5 percent for the initial two years, followed by a mild increase (0.4 percentage points) for the next eight years.[48]

These policy effects are simulated as a deviation from the model's steady state. Obviously, the U.S. economy was in the midst of severe financial turbulence when government funds were injected into financial institutions. Since a log-linearized version of the model is used for the simulation, a simple transfer policy would bring about identical results regardless of where the simulation starts. However, the effects of a debt relief can be different even in a linear model due to the introduction of default risk. The impatient households' decision whether to default or not depends on their outstanding debt. The amount of existing debt at any period during the simulation is not necessarily the same because it is endogenously determined as part of the equilibrium. On top of that, the federal fund rates have remained near zero since 2008:Q4. Taking into account that the federal fund rates are constrained at the zero lower bound (ZLB) can create additional non-linear effects on the economy. The subsequent experiments deal with these two considerations.

I again use the previous crisis experiment for evaluating the effects of the two policies. To measure policy effects I simulate the model with and without a policy shock and compute the differences in simulation outcomes. More precisely, I first compute the model implications only with a series of negative housing preference shocks, and then compute the model implications with an additional policy shock in 2009:Q1 (the twelfth quarter).

Figure 6 depicts the effects of bank recapitalization (solid line) versus household debt relief (dotted line) in the middle of the U.S. housing debt crisis. The left column shows the policy effects without the consideration of the ZLB. The vertical line indicates the period when the policy action is undertaken. The effects of household debt relief on consumption are striking. The debt relief increases consumption more than 3 percentage points relative to the consumption level that would prevail without the policy. In the experiment starting from the steady state, the same policy raises consumption up to around 1.8 percentage points, which is far below 3 percentage points. As shown in panel G of figure 5, charge-off rates on home

[48] Policy effects are computed by just multiplying the simulation outcomes given in figure 4 with the relative magnitude of the policy (4.2).

Figure 6. Effects of Bank Recapitalization versus Houshold Debt Relief Worth 4.2 Percent of Annual GDP in the Middle of the U.S. Housing Debt Crisis

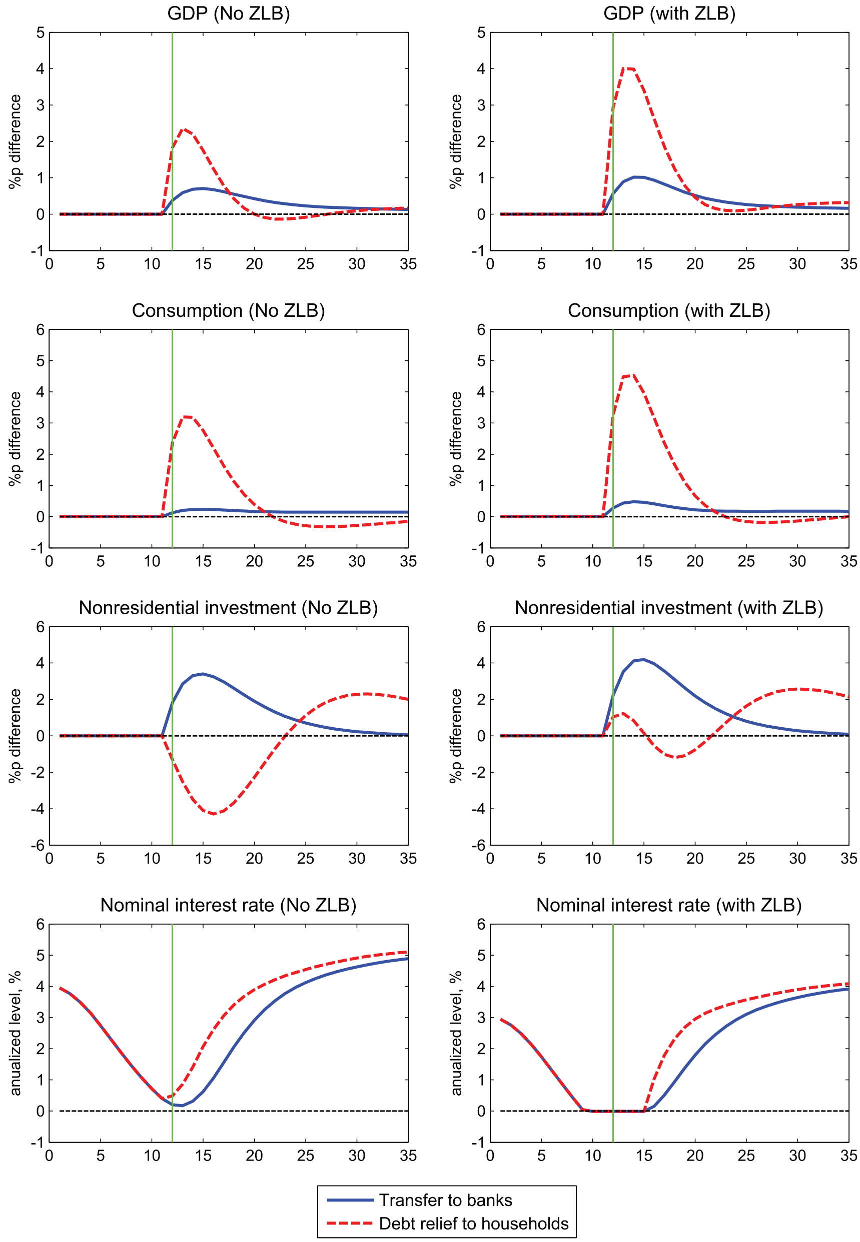

Notes: The simulations start from the steady state. The vertical axis represents the differences in model implications with a policy shock and those without a policy shock. The vertical line in the twelfth quarter indicates the period when a policy action is undertaken. Please refer to the text for more details.

mortgages are much more elevated in 2008:Q4 and thus the lending rate spread of home mortgages is very high, reflecting high default risk. Reducing the households' indebtedness in such an environment, by forgiving part of household debt, improves the lending conditions much more than in normal times. The debt relief causes a decrease in the annualized lending rate spread of home mortgages from 2.0 percentage points to 0.6 percentage points, while the same policy in normal times results in a decline of only 0.8 percentage points. It also brings about further reductions in foreclosure costs. As a result, the aggressively lowered funding costs boost the consumption of credit-constrained households more. A surge in consumption, however, causes higher inflation. Due to the central bank's reaction, interest rates rise more, which in turn leads non-residential investment to decline more than it does in normal times.

Overall, the household debt relief worth 4.2 percent of annual GDP increases GDP by 2.7 percentage points for the first two years, and by 0.7 percentage points for the next eight years. Due to a linear model, the effects of bank recapitalization, a simple transfer to banks, are the same as those in normal times: 1.2 percentage points for the first two years and 1.3 percent for the next eight years. Therefore, the debt relief to credit-constrained households can have much stronger short-run stimulus on output when assuming that the economy is in the middle of a housing debt crisis.

4.2.3 Policy Experiments with the Consideration of the ZLB

I now examine how the consideration of the zero lower bound (ZLB) influences the stimulative effects of both policies. In the experiment with regard to the U.S. housing debt crisis, the simulated nominal interest rate declines close to zero but never hits the ZLB. To increase the likelihood of the nominal rate's hitting the ZLB with few changes of the model dynamics in a linear model, I lower the steady-state annual inflation from 2 percent to 1 percent. I use a toolkit provided by Guerrieri and Iacoviello (2015b) to compute dynamic responses in the presence of the ZLB. The toolkit defines the regime based on whether the constraint binds or not, and uses a piecewise linear perturbation method that basically links the first-order approximation of the model around the same point under each regime. As the dynamics in each regime could depend on the expectation of how long the economy stays in that regime, the simulation outcomes can be highly non-linear.

As in the previous policy experiment, I feed in each policy shock in the twelfth quarter. Without any policy, the nominal interest rate would stay at zero for five quarters more until the sixteenth quarter.[49] Figure 6 presents the simulation results for GDP, consumption, investment, and the nominal interest rate without the ZLB on the left and with the ZLB on the right. First of all, the output effects of both policies are bigger when the nominal interest rate is constrained at zero than in the other case. The peak output effect of the transfer to banks under the ZLB is approximately 1.0 percentage point, which is around 40 percent higher than that without the ZLB (0.7 percentage point). The peak output effect of the debt relief to impatient households under the ZLB is roughly 4.0 percentage points, which is nearly 70 percent higher than that of normal times (2.4 percentage points). Second, each policy helps the economy to escape from a liquidity trap one quarter earlier. The nominal interest rate is constrained at zero for one year until the fifteenth quarter and thereafter it is set according to the interest rate rule. Third, the stimulative effects of the debt relief are magnified under the ZLB. The debt relief is inflationary, so that the real interest rate falls, when the nominal interest rate cannot go below zero. The decrease in real interest rates in turn boosts consumption of both households as well as non-residential investment, and at least weakens crowding-out effects.

Overall, a debt relief worth 4.2 percent of annual GDP increases annual GDP by 5.2 percentage points for the first two years. These output effects are far beyond those of bank recapitalization (1.6 percentage points). For the next eight years, the debt relief would increase GDP by 2.0 percent, whereas the bank capitalization boosts GDP by 1.5 percent. These results suggest that the debt relief policy under the ZLB has far-reaching effects on the economy beyond the periods during which the nominal interest rate is constrained at

[49]The nominal interest rate is endogenously kept at zero for two years from the ninth to the sixteenth quarter. Since the ZLB prevents the nominal interest rate from falling down to the level at which efficient intertemporal allocations can be made, aggregate demand consisting of consumption and non-residential investment decreases further than without the ZLB. The discussion and figures are presented in Yoo (2017).

Table 4. Output Effects of Each One-Time Policy Worth 4.2 Percent of Annual GDP

	Short Run (~2 Years)		Long Run (3–10 Years)	
	Transfer to Banks	Debt Relief to Households	Transfer to Banks	Debt Relief to Households
Normal Times	1.2	1.5	1.3	0.4
Housing Debt Crisis	1.2	2.7	1.3	0.7
Housing Debt Crisis under the ZLB	1.6	5.2	1.5	2.0

Note: Output effects are measured relative to the baseline annual GDP.

Figure 7. Ten-Years Cumulative Output Multipliers and The Degree of Monetary Accommodation

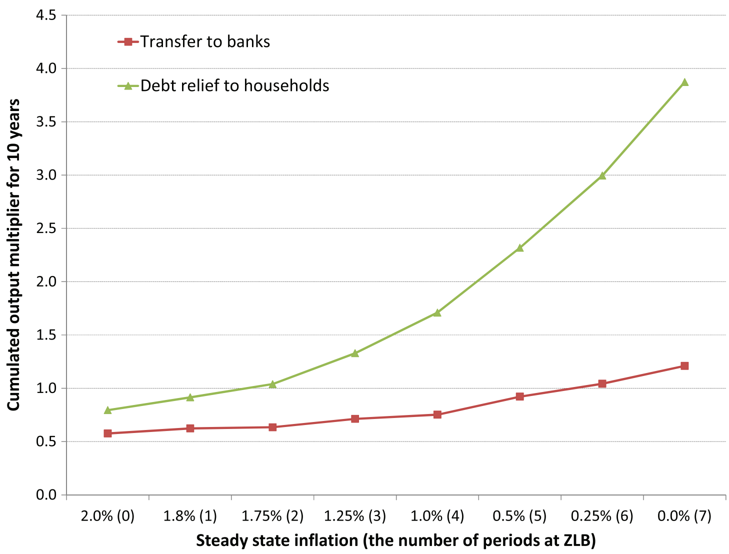

zero. Table 4 summarizes the comparison of counterfactual output effects in three circumstances.

I also conduct the same policy simulations with a different steady-state inflation to show how the degree of monetary accommodation affects the stimulative effects of both policies. Figure 7 presents the cumulative output multipliers of each policy for ten years—the line with squares for the transfer to banks and the line with triangles for the debt relief to households—at each steady-state inflation. Following Uhlig (2010), the cumulative multiplier is defined as a ratio whose numerator equals the present values of policy-induced output effects relative to the steady state for ten years, and whose denominator equals the size of the policy measure, which in this case is 4.2 percent of annual steady-state GDP. I assume zero discount rates, namely, making a simple sum of policy-induced output effects.[50] The number in the parentheses on the x-axis indicates how many quarters the nominal interest rates are stuck at zero from the period of the policy intervention. As can be seen, the longer the policy rates are constrained at the ZLB, the higher the cumulative output multipliers are for both policies. The cumulative multiplier of capital injection policies is 0.6 when the ZLB is disregarded, but it doubles when the policy rate stays at zero for seven quarters. More interestingly, in the case of debt relief policies, the cumulative multiplier accelerates as the number of periods that the economy is constrained at the ZLB increases. The cumulative multiplier when the policy rate stays at zero for seven quarters is nearly 4.0, which is nearly five times as big as that without the ZLB (0.8).

5. Conclusion

In this paper, I build a DSGE model having major macrofinancial linkages identified in the literature with a particular focus on the role of housing finance. I use the model to analyze the macroeconomic effects of a collapse in house prices observed in the recent U.S. housing debt crisis. Then such a crisis environment is used as a laboratory for assessing the relative effectiveness of a policy to inject capital into banks versus a policy to relieve households of mortgage debt. The magnitude and timing of each policy are chosen to be comparable to those of financial-sector support programs that the U.S. government conducted during the crisis.

The policy simulation results suggest that policies to support the financial sector helped to ease the severity of the recession in

[50]This is equivalent to the assumption that the discount rates are the same as the trend growth rates.

the aftermath of the housing debt crisis. The more important implication is, however, that debt relief to households of the same amount of funds would have been much more effective to stimulate the economy, given that the household sector was highly leveraged and the nominal interest rates were constrained at the ZLB. Therefore, this paper supports the view of Mian and Sufi (2014) and others, who argue for a more aggressive reduction of household debt.[51]

Despite the detailed model structure, several issues need to be addressed in order to quantify the effects of household debt relief more precisely. First, it would take some time for government agencies to undertake the necessary steps to carry out a policy to reduce household debt, because such a policy involves a lot of parties unlike the financial-sector support programs. In my policy experiments, however, the debt relief is modeled as a one-time policy intervention to facilitate a direct comparison with the capital injection to banks. As shown in Cogan et al. (2010), the significant implementation lag of fiscal policy can reduce its stimulative effects. Meanwhile, the anticipation of the scheduled debt relief policy can ameliorate the current credit condition by loosening the participation constraint of banks. In the end, it is a quantitative issue as to which one dominates the other. Second, while the model partly accounts for sluggish mortgage dynamics by introducing loan adjustment costs, the introduction of explicitly defined long-term mortgage contracts would be more desirable. The recent papers by Elenev, Landvoigt, and Nieuwerburgh (2015) and Gelain, Lansing, and Natvik (2015) deal with an interesting modeling setup for long-term mortgage contracts. Third, the default decision could be modeled in more detail. For example, by introducing a direct penalty or cost on defaulted agents, an access to credit market would be restricted for a stochastic length of periods once borrowers declare default (see Chatterjee and Eyigungor 2015).

Besides the evaluation of the post-crisis government policies, the model can be used or extended for analyzing other interesting macroeconomic policies. Areas for future research include the design of

[51]Note that throughout the policy simulations, I implicitly assume that the main driver of the recent housing debt crisis is a shock on house prices. If other shocks, such as a shock originating from the financial sector are considered to be dominant, the relative effectiveness of the policies can be different.

optimal monetary policy rules with financial variables (as in Cúrdia and Woodford 2010, Gilchrist and Zakrajsek 2012, and Hirakata, Sudo, and Ueda 2013) and the optimal bank capital regulation (as in Clerc et al. 2015 and Begenau 2016).

References

Begenau, J. 2016. "Capital Requirements, Risk Choice, and Liquidity Provision in a Business Cycle Model." Working Paper.

Benes, J., and M. Kumhof. 2011. "Risky Bank Lending and Optimal Capital Adequacy Regulation." IMF Working Paper No. 11/130.

Bernanke, B. S. 2015. *The Courage to Act: A Memoir of a Crisis and Its Aftermath.* W. W. Norton & Company.

Bernanke, B. S., M. Gertler, and S. Gilchrist. 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." In *Handbook of Macroeconomics,* Vol. 1, Part C, ed. J. B. Taylor and M. Woodford, 1341–93 (chapter 21). Elsevier.

Carlstrom, C. T., and T. S. Fuerst. 1997. "Agency Costs, Net Worth, and Business Fluctuations: A Computable General Equilibrium Analysis." *American Economic Review* 87 (5): 893–910.

Chatterjee, S., and B. Eyigungor. 2015. "A Quantitative Analysis of the U.S. Housing and Mortgage Markets and the Foreclosure Crisis." *Review of Economic Dynamics* 18 (2): 165–84.

Christiano, L., and D. Ikeda. 2013. "Leverage Restrictions in a Business Cycle Model." NBER Working Paper No. 18688.

Christiano, L. J., R. Motto, and M. Rostagno. 2014. "Risk Shocks." *American Economic Review* 104 (1): 27–65.

Clerc, L., A. Derviz, C. Mendicino, S. Moyen, K. Nikolov, L. Stracca, J. Suarez, and A. P. Vardoulakis. 2015. "Capital Regulation in a Macroeconomic Model with Three Layers of Default." *International Journal of Central Banking* 11 (3): 9–63.

Coenen, G., C. J. Erceg, C. Freedman, D. Furceri, M. Kumhof, R. Lalonde, D. Laxton, J. Linde, A. Mourougane, D. Muir, S. Mursula, C. de Resende, J. Roberts, W. Roeger, S. Snudden, M. Trabandt, and J. in't Veld. 2012. "Effects of Fiscal Stimulus in Structural Models." *American Economic Journal: Macroeconomics* 4 (1): 22–68.

Cogan, J., T. Cwik, J. Taylor, and V. Wieland. 2010. "New Keynesian versus Old Keynesian Government Spending Multipliers." *Journal of Economic Dynamics and Control* 34 (3): 281–95.

Covas, F., and J. C. Driscoll. 2014. "Bank Liquidity and Capital Regulation in General Equilibrium." Finance and Economics Discussion Series No. 2014-85, Board of Governors of the Federal Reserve System.

Cúrdia, V., and M. Woodford. 2010. "Credit Spreads and Monetary Policy." *Journal of Money, Credit and Banking* 42 (s1): 3–35.

Dogra, K. 2014. "Liquidity Traps, Debt Relief, and Macroprudential Policy: A Mechanism Design Approach." Working Paper.

Duca, J. V., J. Muellbauer, and A. Murphy. 2011. "House Prices and Credit Constraints: Making Sense of the US Experience." *Economic Journal* 121 (552): 533–51.

Eggertsson, G. B., and P. Krugman. 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *Quarterly Journal of Economics* 127 (3): 1469–1513.

Elenev, V., T. Landvoigt, and S. V. Nieuwerburgh. 2015. "Phasing Out the GSEs." Working Paper.

Forlati, C., and L. Lambertini. 2011. "Risky Mortgages in a DSGE Model." *International Journal of Central Banking* 7 (1): 285–335.

Gale, D., and M. Hellwig. 1985. "Incentive-Compatible Debt Contracts: The One-Period Problem." *Review of Economic Studies* 52 (4): 647–63.

Geanakoplos, J. 2010. "The Leverage Cycle." In *NBER Macroeconomics Annual 2009*, Vol. 24, ed. D. Acemoglu, K. Rogoff, and M. Woodford, 1–65 (chapter 1). University of Chicago Press.

Geithner, T. F. 2014. *Stress Test: Reflections on Financial Crises.* Crown Publishers.

Gelain, P., K. J. Lansing, and G. J. Natvik. 2015. "Leaning against the Credit Cycle." Working Paper No. 04/2015, Norges Bank.

Gerali, A., S. Neri, L. Sessa, and F. M. Signoretti. 2010. "Credit and Banking in a DSGE Model of the Euro Area." *Journal of Money, Credit and Banking* 42 (s1): 107–41.

Gertler, M., and P. Karadi. 2011. "A Model of Unconventional Monetary Policy." *Journal of Monetary Economics* 58 (1): 17–34.

Gilchrist, S., and E. Zakrajsek. 2012. "Credit Supply Shocks and Economic Activity in a Financial Accelerator Model." In

Rethinking the Financial Crisis, ed. A. S. Blinder, A. W. Lo, and R. M. Solow, 37–72 (chapter 3). Russel Sage Foundation.

Guerrieri, L., and M Iacoviello. 2015a. "Collateral Constraints and Macroeconomic Asymmetries." Working Paper.

———. 2015b. "OccBin: A Toolkit for Solving Dynamic Models with Occasionally Binding Constraints Easily." *Journal of Monetary Economics* 70 (March): 22–38.

Guerrieri, L., M. Iacoviello, F. Covas, J. C. Driscoll, M. Jahan-Parvar, M. Kiley, A. Queralto, and J. Sim. 2015. "Macroeconomic Effects of Banking Sector Losses across Structural Models." Finance and Economics Discussion Series No. 2015-44, Board of Governors of the Federal Reserve System.

Hirakata, N., N. Sudo, and K. Ueda. 2013. "Capital Injection, Monetary Policy, and Financial Accelerators." *International Journal of Central Banking* 9 (2): 101–45.

Iacoviello, M. 2005. "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle." *American Economic Review* 95 (3): 739–64.

———. 2015. "Financial Business Cycles." *Review of Economic Dynamics* 18 (1): 140–64.

Iacoviello, M., and S. Neri. 2010. "Housing Market Spillovers: Evidence from an Estimated DSGE Model." *American Economic Journal: Macroeconomics* 2 (2): 125–64.

Jeske, K., D. Krueger, and K. Mitman. 2013. "Housing, Mortgage Bailout Guarantees and the Macro Economy." *Journal of Monetary Economics* 60 (8): 917–35.

Jordá, O., M. Schularick, and A. M. Taylor. 2014. "The Great Mortgaging: Housing Finance, Crises, and Business Cycles." NBER Working Paper No. 20501.

Justiniano, A., G. E. Primiceri, and A. Tambalotti. 2015. "Household Leveraging and Deleveraging." *Review of Economic Dynamics* 18 (1): 3–20 (Special Issue: Money, Credit, and Financial Frictions).

Kaplan, G., and G. L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82 (4): 1199–1239.

Kiley, M. T., and J. W. Sim. 2014. "Bank Capital and the Macroeconomy: Policy Considerations." *Journal of Economic Dynamics and Control* 43 (C): 175–98.

———. 2015. "Optimal Monetary and Macroprudential Policies: Gains and Pitfalls in a Model of Financial Intermediation." Finance and Economics Discussion Series No. 2015-78, Board of Governors of the Federal Reserve System.

Kiyotaki, N., and J. Moore. 1997. "Credit Cycles." *Journal of Political Economy* 105 (2): 211–48.

Kollmann, R., M. Ratto, W. Roeger, and J. in't Veld. 2013. "Fiscal Policy, Banks and the Financial Crisis." *Journal of Economic Dynamics and Control* 37 (2): 387–403 (Special Issue: Fiscal Policy in the Aftermath of the Financial Crisis).

Landvoigt, T. 2014. "The Allocation and Pricing of Mortgage Risk." Working Paper.

Liu, Z., P. Wang, and T. Zha. 2013. "Land-Price Dynamics and Macroeconomic Fluctuations." *Econometrica* 81 (3): 1147–84.

Mian, A., K. Rao, and A. Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128 (4): 1687–1726.

Mian, A., and A. Sufi. 2014. *House of Debt: How They (and You) Caused the Great Recession, and How We Can Prevent It from Happening Again*. University of Chicago Press.

Mian, A., A. Sufi, and F. Trebbi. 2015. "Foreclosures, House Prices, and the Real Economy." *Journal of Finance* 70 (6): 2587–2634.

Nakamura, E., and J. Steinsson. 2014. "Fiscal Stimulus in a Monetary Union: Evidence from US Regions." *American Economic Review* 104 (3): 753–92.

Quint, D., and P. Rabanal. 2014. "Monetary and Macroprudential Policy in an Estimated DSGE Model of the Euro Area." *International Journal of Central Banking* 10 (2): 169–236.

Reinhart, C., and K. Rogoff. 2009. *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press.

Shiller, R. 2012. "Reviving Real Estate Requires Collective Action." *New York Times*, June 23.

Smets, F., and R. Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586–606.

Stiglitz, J. 2010. "Foreclosures and Banks' Debt to Society." *The Guardian*, November 5.

Summers, L. 2014. "Lawrence Summers on 'House of Debt': Did the Response to the Financial Crisis Focus Too Much on Banks while Neglecting Over-indebted Homeowners?" *Financial Times*, June 6.

Uhlig, H. 2010. "Some Fiscal Calculus." *American Economic Review* 100 (2): 30–34.

van der Kwaak, C., and S. van Wijnbergen. 2014. "Financial Fragility, Sovereign Default Risk and the Limits to Commercial Bank Bail-outs." *Journal of Economic Dynamics and Control* 43 (C): 218–40.

Walentin, K. 2014. "Business Cycle Implications of Mortgage Spreads." *Journal of Monetary Economics* 67 (C): 62–77.

Whelan, K. 2002. "A Guide to U.S. Chain Aggregated NIPA Data." *Review of Income and Wealth* 48 (2): 217–33.

Wieland, V., E. Afanasyeva, M. Kuete, and J. Yoo. 2016. "New Methods for Macro-Financial Model Comparison and Policy Analysis. In *Handbook of Macroeconomics*, Vol. 2, ed. J. B. Taylor and H. Uhlig, 1241–1319 (chapter 15). Elsevier.

Yoo, J. 2017. "Capital Injection to Banks versus Debt Relief to Households." IMFS Working Paper No. 111.

Discussion of "Capital Injection to Banks versus Debt Relief to Households"*

Atif Mian
Princeton University and NBER

Jinhyuk Yoo asks an important and interesting question in this paper: if policymakers have to choose between providing debt relief to households versus injecting new capital into banks, how should they evaluate the choice? The paper addresses this question through the lens of a DSGE model, and makes the case that household debt relief is more potent than capital injection for banks as a stimulus measure in the midst of a recession drive by household debt like the one the United States experienced in 2008. Amir Sufi and I also came to a similar conclusion in our 2014 book *House of Debt*. In my comments below, I will first describe the intuition behind the key result in Yoo's work. I then discuss the empirical relevance of the model's assumption and implications in light of the increasing body of work that has come about since the great recession. Finally, I discuss a conceptual extension of the model that Yoo presents to highlight some important questions for further inquiry.

1. Basic Intuition

The paper builds on a standard New Keynesian DSGE model with the typical ingredients: (i) a Euler equation that connects consumption today to consumption tomorrow through the interest rate and expected inflation, (ii) a monetary authority that follows the Taylor principle subject to a zero lower bound (ZLB) constraint, and (iii) a Phillips curve for the supply side with a positive relationship between output and inflation. It is well known that the standard DSGE model with these three ingredients does not have a direct role for financial policy, either for household debt relief or for bank capital injection.

*Author contact: atif@princeton.edu.

Therefore, to make the model relevant for the question at hand, Yoo adds two additional ingredients.

First, there is heterogeneity across households, with patient households lending to impatient households. The impatient households have higher marginal propensity to consume (MPC) but face a leverage constraint that limits how much they can borrow from patient households. Second, there is a banking sector that funds both households and entrepreneurs. However, the banking sector faces a leverage constraint that makes the net worth of the banking sector relevant for aggregate credit supply.

With the addition of the above ingredients, the models can address the key policy question in a non-trivial way. Since impatient households have higher marginal propensity to consume, tighter constraint on their spending can reduce aggregate demand if the economy is unable to stimulate additional demand due to interest rates stuck at the zero lower bound. In such a situation, "redistributive policy" can expand total output by expanding demand. For example, a transfer from the patient to the impatient households in the form of "debt relief" raises aggregate demand due to the above-mentioned higher MPC for impatient households in the presence of zero lower bound. In fact, the paper points out a second important advantage of expanding demand, namely its positive impact on inflation. When the nominal interest rate is stuck at zero, any increase in expected inflation is helpful, as it lowers the real rate and brings the economy closer to its natural rate of interest.

The model weighs the positive gains from household debt relief due to the above forces against the potential gain from injecting new capital into banks. For example, instead of transferring funds from patient households to impatient households in the form of debt relief, we can transfer funds from patient households to banks in the form of capital injection. An increase in bank capital helps to lower spread. This raises investment and also enables impatient households to borrow more at lower rates, thus boosting demand. However, the expansion in investment credit to entrepreneurs expands supply and hence is fundamentally deflationary in scope. This is the key effect that makes bank capital intervention relatively less attractive in a zero lower bound environment: the economy needs higher inflation expectations, and a deflationary shock is not helpful in achieving that goal.

The model thus operates by translating household debt relief into an outward expansion of the aggregate demand curve, which is inflationary in nature. On the other hand, bank capital injection translates into an expansion of the aggregate supply curve, which is deflationary in nature. *This asymmetric effect is key to understanding the main result of the paper.* When the source of the recession is a fall in aggregate demand due to household debt overhang, the economy is in a deflationary cycle with interest rate at the zero lower bound. In such a scenario, policy that is inflationary in nature has a stronger multiplier—in the paper's calibration, the effect is three times as large in the short run. This effect is specific to a ZLB world and does not exist otherwise due to the Taylor rule counteracting any effect on inflation.

2. Empirical Relevance

The paper builds a nice intuitive model to make the case for household debt relief in a demand-constrained economy. However, how realistic are the model's assumptions and quantitative predictions—especially with reference to the 2008 financial and economic crisis? I discuss two of the model's assumptions, (i) regarding heterogeneity in consumer preferences, and (ii) regarding assumed deviation from steady state. I then discuss the empirical relevance for the key prediction that household debt relief would lead to strong consumption response from indebted households.

The first key assumption of the model is a persistent and quantitatively important difference in marginal propensity to consume across households in the economy. This, by now, should be an uncontroversial assumption. There are a number of careful empirical studies that document strong differences in marginal propensity to consume and marginal propensity to borrow out of financial shock. Mian, Rao, and Sufi (2013) estimate large differences in marginal propensity to consume out of wealth shocks, with poorer households and households with high leverage having much stronger MPC out of wealth. Mian and Sufi (2011) show that homeowners with lower credit scores and homeowners with liquidity constraints borrow significantly more for a dollar increase in house price relative to the rest of the population. Agarwal et al. (2015) document that individuals with low credit scores respond much more strongly to an increase in

Figure 1. Consumption-to-Income Ratio over Time

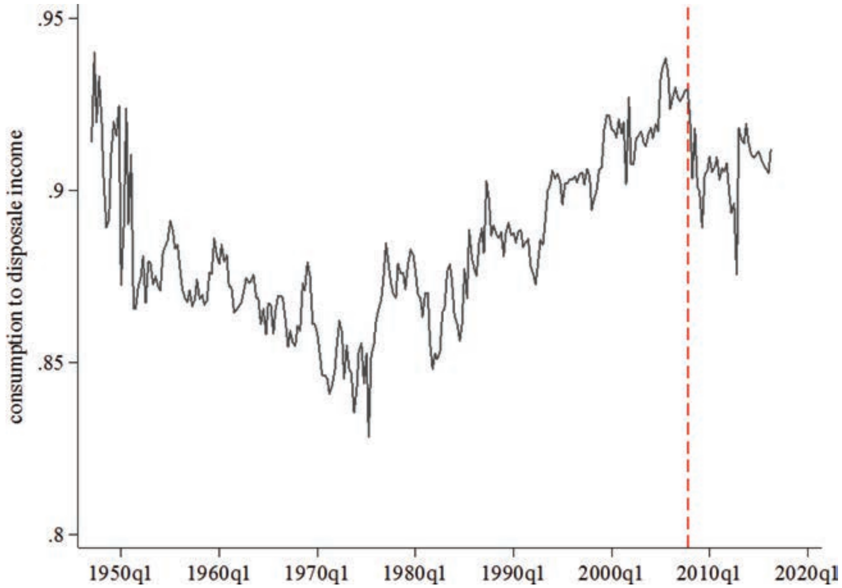

credit limit, for borrowing as well as spending. Parker et al. (2013) estimate that lower-income households and liquidity-constrained households respond much more aggressively to fiscal stimulus checks.

A second assumption of the model is that the economy is in a steady state when the adverse shock hits. This is a natural and reasonable assumption. However, to the extent we want to apply the model to the 2008 recession, it is important to highlight that the economy on the eve of the crisis was not in "steady state." In particular, credit dynamics leading up to 2007 may themselves reflect deviation from the norm. This can be seen in figure 1, which plots the ratio of consumption to disposable income over time. The data for this figure come from Pistaferri (2016).

The consumption-to-income ratio on the eve of the crisis was at a post-war high, having risen significantly in the preceding two decades. At the same time, household leverage was also at an all-time high, having undergone a rapid acceleration in the decade before. The two figures are likely related. As has been shown in prior work, many homeowners borrowed aggressively against the rising value of their homes in the 2000s. Such borrowing is likely to provide a short-term boost to consumption. There is thus a close empirical relationship between the peak in the ratio of consumption to disposable income in 2006 and the peak in household leverage to GDP around the same time. If the consumption-to-income ratio tends to revert to its longer-term mean, then we would naturally expect consumption growth to slow down relative to income post-2008 crisis.

Thus there may be additional forces in the background that make it difficult to sustain the level of consumption relative to disposable income seen in 2006. For example, if the higher consumption-to-income level was sustained due to the high credit growth during the 2000s, the ratio will naturally need to come down as the economy settles into a more normal flow of credit creation. Such an adjustment would create an additional headwind for consumption growth going forward. I suspect that if we take such additional forces into account, the positive impact of debt relief due to its inflationary pressure may be even further enhanced. We could benefit from such an analysis in the paper.

A key prediction of the model's framework is that debt relief leads to an increase in consumption for the indebted consumer. This follows naturally from the assumption of heterogeneity in MPC already discussed above. However, from an empirical standpoint, heterogeneity in MPC with respect to income or other financial shocks does not necessarily mean that households respond symmetrically to a write-down in debt balance as well. There is an increasing empirical literature that helps illuminate this issue.

Di Maggio, Kermani, and Ramcharan (2014) and Keys et al. (2014) design a natural experiment that exploits predetermined differences in the timing of when mortgage interest rates get set to a lower amount as monetary policy relaxes in the aftermath of the 2007–08 crisis. The authors show that households that receive a lower mortgage interest rate are less likely to default and also use that net cash saving to boost spending. However, cash flow constraints are important in understanding these results. Ganong and Noel (2016) use another natural experiment that explores debt write-down for "underwater" homeowners. The authors show that households experiencing a write-down on principal payment alone, without a reduction in cash payment due to the lender, does not result in a spending increase. Thus the combined evidence suggests that interest burden relief is essential to benefit in terms of spending gain. Otherwise, individuals may be too cash flow constrained to benefit from a pure principal write-down. These details are important when quantitatively assessing the model's implications.

3. Model Extension

An important conceptual insight in the paper is that government intervention in financial markets can affect either the aggregate demand or the aggregate supply side of the economy depending on the nature of government intervention. This difference makes the nature of government intervention important, since a policy that expands demand is inflationary and hence more helpful in situations that suffer from liquidity trap. A related insight is also present in Bahadir and Gumus (2016), who explicitly model the household and firm sectors separately.

A policy shock in the Bahadir and Gumus (2016) paper is inflationary if it primarily expands the household sector and hence aggregate demand. However, a similar shock affecting the firm sector is deflationary, as its primary impact is on expanding aggregate supply. An advantage of the Bahadir and Gumus (2016) approach is that one can generate further empirical predictions to test whether a particular policy shock largely worked through the demand side or the supply side.

Mian, Sufi, and Verner (2017) introduce a tradable and non-tradable sector in the economy that otherwise looks similar to the Bahadir and Gumus (2016) approach. We show in that paper that credit supply shocks that boost household demand will tend to raise employment in the non-tradable sector relative to the tradable sector. At the same time, higher demand will push up the prices of local non-tradable goods. In contrast, credit supply shocks that expand tradable or non-tradable firms' labor productivity do not make this joint prediction.

We can thus use data on exployment and prices at the tradable and non-tradable sector level to understand whether a given shock operated through the firm side of households. The empirical results in Mian, Sufi, and Verner (2017) suggest that credit shocks during the 1980s primarily worked through the household sector by affecting aggregate demand.

The broader comment I want to make here is that the author can look at tradable and non-tradable sectors separately to tease apart the true relevance of the mechanisms highlighted in the paper.

4. Conclusion

Overall, this is an excellent paper that makes the important point that policies focused on debt relief, or lower mortgage payment, have a more stimulative effect than policies focused on bank capital injection when the source of crisis is insufficient demand. The key "multiplier" is driven by the inflationary impact of providing debt relief to highly indebted households. Future extensions of the paper can benefit from incorporating more of the empirical insights in the literature, and can enrich theory further by separating the tradable and non-tradable sectors. Doing so delivers a new set of predictions that can then be taken to data in subsequent work.

References

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel. 2015. "Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?" NBER Working Paper No. 21567.

Bahadir, B., and I. Gumus. 2016. "Credit Decomposition and Business Cycles in Emerging Market Economies." *Journal of International Economics* 103: 250–62.

Di Maggio, M., A. Kermani, and R. Ramcharan. 2014. "Monetary Policy Pass-Through: Household Consumption and Voluntary Deleveraging."

Ganong, P., and P. Noel. 2016. "The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession." Working Paper.

Keys, B. J., T. Piskorski, A. Seru, and V. Yao. 2014. "Mortgage Rates, Household Balance Sheets, and the Real Economy." NBER Working Paper No. 20561.

Mian, A., K. Rao, and A. Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128 (4): 1687–1726.

Mian, A., and A. Sufi. 2011. "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis." *American Economic Review* 101 (5): 2132–56.

Mian, A. R., A. Sufi, and E. Verner. 2017. "Credit Supply and Business Cycle Amplification: Evidence from Banking Deregulation in the 1980s."

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103 (6): 2530–53.

Pistaferri, L. 2016. "Why Has Consumption Remained Moderate after the Great Recession?" Working Paper.

Policies for Crises Prevention and Management*

Arvind Krishnamurthy
Stanford University, Graduate School of Business and National Bureau of Economic Research

I discuss the role of debt in financial crises and macroprudential tools to manage debt. A principal conclusion is that changes in mortgage contract designs, an ex ante measure aimed at the household sector, can deliver significant welfare benefits.

JEL Codes: G01, G21, G28, E44, E58.

The papers by Griss, Krogstrup, and Schumacher; Yoo; and Bahaj and Foulis (all in this issue) study macroprudential policy tools to improve financial stability. The policies include ex post tools in response to financial crises and ex ante tools to reduce the likelihood and severity of a crisis. Griss, Krogstrup, and Schumacher examine an ex post tool, studying the effective lower bound on monetary policy. Yoo also studies ex post policies, seeking to quantify the macroeconomic benefits of household debt relief relative to bank capital injections during a crisis. Bahaj and Foulis examine ex ante macroprudential policies and the role of uncertainty regarding the efficacy of macroprudential policy tools in designing optimal policy.

My comments highlight the role of debt in financial crises and macroprudential tools to manage debt. I discuss ex ante and ex post tools, uncertainty over policy outcomes, and the distinction between debt relief to the household sector and the financial sector. A principal conclusion is that changes in mortgage contract designs, an ex ante measure aimed at the household sector, can deliver significant welfare benefits.

*I thank Jonathan Wallen for research assistance. Author e-mail: a-krishnamurthy@stanford.edu.

1. 1998 versus 2008

Figure 1 graphs the evolution of the Gilchrist-Zakrajsek spread (GZ spread, from Gilchrist and Zakrajsek 2012) and the three-month commercial paper to three-month Treasury-bill spread over two financial crises. The top panel spans January 1998 to June 1999, including the hedge fund crisis of the fall of 1998. The bottom panel graphs these spreads from January 2007 to July 2008, a period that covers part of the 2007–09 financial crisis.

In both events, spreads rise sharply, reflecting a flight to liquidity and safety as well as increased concerns regarding a recession and default risks. But in October 1998, the spreads begin to fall, coincident with the Federal Reserve's interventions (see Scholes 2000). In the 2007–09 crisis, the spreads continue to rise through this entire period, peaking in the fall of 2008. What distinguishes these events? Why in 1998 did the financial disruption as reflected in asset markets dissipate with no real U.S. macroeconomic consequences?

2. Debt Is the Source of Vulnerability

Debt is the economic channel connecting financial distress with the real economic outcomes. There is a significant debt buildup prior to 2007 and virtually no debt buildup prior to 1998. Figure 2 graphs the aggregate quantity of short-term debt liabilities of the financial sector, normalized by GDP, over the period from 1990 to 2014. The data is taken from Krishnamurthy and Vissing-Jorgensen (2015), who discuss the data construction. While the data in figure 2 measure bank debt, there is also a dramatic rise in household indebtedness ahead of the 2008 event (Mian and Sufi 2015).

This comparison speaks to a more general pattern. Jordà, Schularick, and Taylor (2013) document from many historical episodes of crises around the world that sharp increases in the credit-to-GDP ratio (primarily bank credit outstanding) increase the likelihood and severity of a financial crisis. Krishnamurthy and Muir (2016) examine the behavior of output, credit, and credit spreads in an international panel of crises episodes, covering fourteen countries and going back to 1869. In this sample, they trace out the path of GDP declines following a large increase in spreads. Conditioning only on a large increase in spreads, they show that expected output declines. This

Figure 1. Crises and Spreads

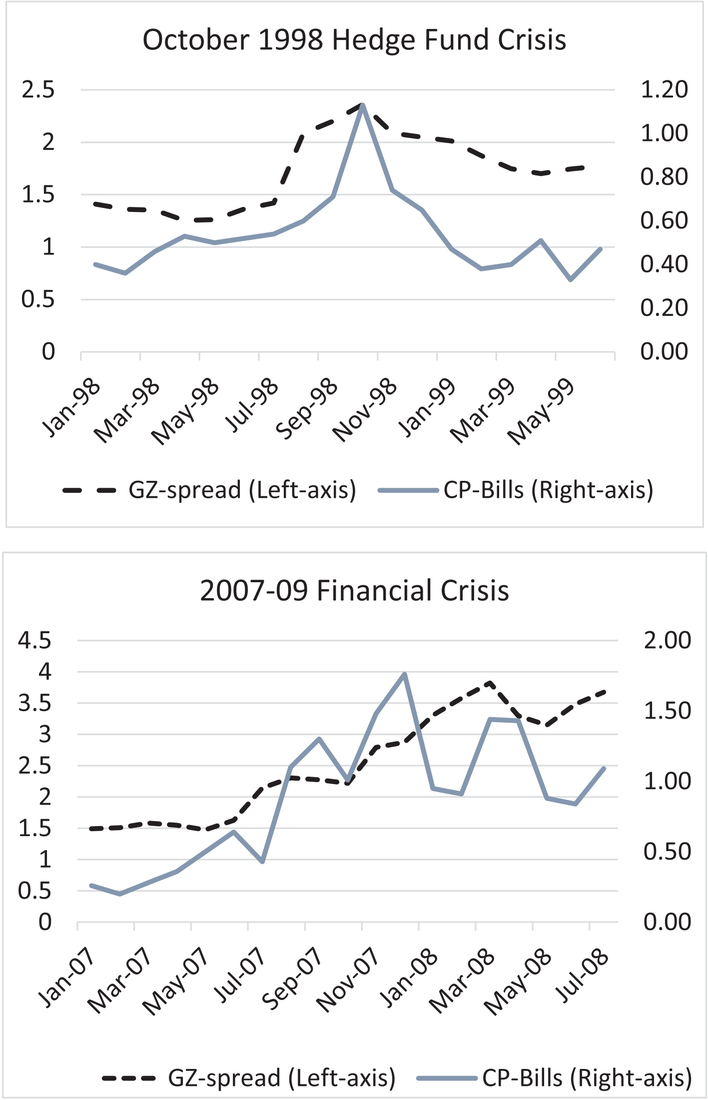

Sources: Simon Gilchrist's webpage and the Federal Reserve.

Figure 2. Short-Term Bank Debt

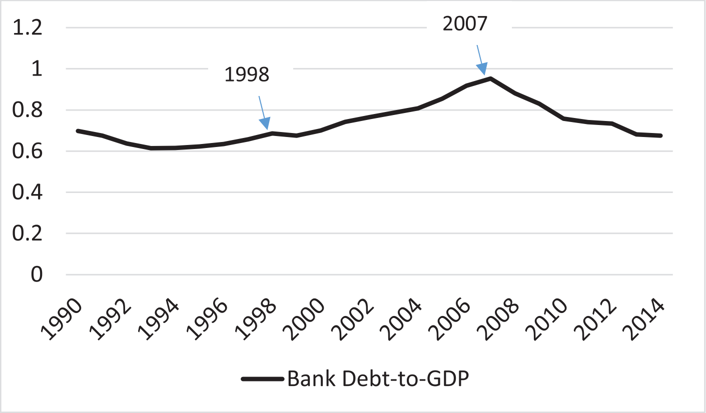

Source: Krishnamurthy and Vissing-Jorgensen (2015).

finding confirms results from the literature on the forecasting power of credit spreads for economic activity (see Gilchrist and Zakrajsek 2012 for U.S. evidence). Conditioning on both a large increase in spreads and high debt levels, expected output declines are much more severe, roughly double that with mean debt levels. In short, the 1998 versus 2008 comparison embodies a general principle that debt is the vulnerability factor.

While there is a growing literature confirming these findings on the role of debt, many important questions remain. Is the vulnerability due to all maturities of debt, or primarily short-term debt as suggested by the banking literature? Does the collateral backing the debt (e.g., real estate versus cars) matter? Pertinent to the analysis of Yoo, the literature has also not quantified the relative importance of non-financial-sector indebtedness versus financial-sector indebtedness.

Managing and preventing crises is about managing the buildup of debt pre-crisis and offsetting the effect of debt overhang in a crisis. This is the post-crisis consensus as reflected in regulatory changes around the world. I next describe the toolkit that governments have used to offset debt overhang and the regulatory reforms designed to prevent debt buildup.

3. Ex Post Policies to Offset the Drag from High Debt Levels

Governments around the world employ similar tools to deal with high debt levels in a crisis. Central banks provide liquidity to the financial sector to minimize funding disruptions and forestall rollover risk. The discount window in all its incarnations was active in the financial crisis both in the United States and around the world. Policy interventions to recapitalize the financial sector with equity sought to reduce the negative effects of debt on the financial sector's activity. Such legislative policy faced a higher hurdle than central bank liquidity provision. The challenges in passing and then implementing the Troubled Asset Relief Program (TARP) have been well chronicled by many of the key actors in the U.S. financial crisis.

Relative to the institutions and policies to manage financial-sector debt, policies that targeted household debt have been far harder to implement. The Federal Reserve's purchases of mortgage-backed securities reduced mortgage rates, but only borrowers with low loan-to-value (LTV) ratios could refinance their mortgages to gain the payment relief from these policies. The Home Affordable Modification Program (HAMP), which was introduced in 2009, aimed to incentivize lenders to renegotiate residential mortgages. But it reached only one-third of eligible borrowers. Moreover, the incidence of renegotiation was driven by intermediary, rather than borrower-specific, factors. In short, given banks' own financial stresses, only some banks rolled out HAMP. The government's other main mortgage restructuring program, the Home Affordable Refinance Program (HARP), also had mixed success. While HARP was initiated in 2009, takeup was low until about 2012, when the program was modified to include high-LTV loans. Moreover, as with HAMP, intermediary-specific factors dampened the impact of the program. These points are documented in two excellent papers by Agarwal et al. (2015, 2016).

Ex post policies aimed at the financial-sector debt problem were much easier to implement than those targeting the household-sector debt problem. Next, consider ex ante policies.

4. Ex Ante Policies

Among both regulators and academics, the post-crisis consensus is that the financial sector should be funded with more equity and less debt. Central to this new financial system are greater capital requirements, tighter leverage requirements, and increased liquidity requirements (see Duffie 2016).

There has been much less done on the household debt front. The Qualified Mortgage rule in the United States reduces payment-to-income ratios on household mortgage debt. Around the world, some financial regulators monitor household LTVs with the goal of restricting high-LTV loans if such loans carry systemic risk.

Banks can satisfy higher capital requirements by raising external equity capital from investors. However, since households do not raise equity, macroprudential policies for household debt tend to constrain credit access. For instance, LTV restrictions constrain the total funding available to households, decreasing household expenditures.

Innovative mortgage designs that are more equity-like are an alternative solution to reduce the negative effects of household debt, with smaller effects on credit access. But the U.S. mortgage market is still dominated by a standard fixed-rate mortgage.

Surprisingly—because they are likely not needed—innovative security designs have been used by the financial sector to avoid raising common equity. Since the crisis, banks have issued contingent convertible securities as well as subordinated debt. In principle, such securities will convert from debt to equity in the event of distress of a financial firm, and thus increase the total loss-absorbing capacity of a bank. However, there are questions about the welfare benefits from these novel contract designs. Admati and Hellwig (2013) have argued, persuasively in my view, that regulators would better achieve their financial stability objectives by focusing purely on common equity rather than chasing new contract designs.

Innovative mortgage contracts, in contrast, can deliver significant macroeconomic benefits. To illustrate, consider the experiences of households with adjustable-rate mortgages (ARMs) during the financial crisis. Fuster and Willen (2013) track a set of borrowers with ARMs over the period from 2009 onwards. Since interest rates tend to fall during recessions, mortgage payments of borrowers fall considerably when their ARMs reset. Fuster and Willen compare different cohorts of ARM borrowers whose rates reset at slightly different calendar times. They show that default and delinquency falls dramatically at the time of the reset for the borrowers receiving the reset compared with the control group of borrowers who receive the reset later. Di Maggio, Kermani, and Ramcharan (2016) document beneficial effects on consumption and debt payments for ARM borrowers with rate resets. These studies reveal the scope for mortgage design, and how alternative designs can beneficially reduce the negative effects of high mortgage debt.

The ARM example is a case of a mortgage design that involves state contingency in payments relative to the standard non-state-contingent fixed-rate mortgage. Mortgage payments were lower in the crisis, since interest rates were positively correlated with household income. There are other forms of state contingency that may also be beneficial. Eberly and Krishnamurthy (2014) propose a mortgage with features of both ARMs and fixed-rate mortgages. Their mortgage is a fixed-rate mortgage with the option to convert to an ARM, or prepay. This small variant on the standard design retains state contingency in payments through the option to convert to an ARM. Guren, Krishnamurthy, and McQuade (2016) study this and other mortgage designs in a general equilibrium model that quantifies the utility to benefits to households of more state contingency in mortgage contracts. Finally, Mian and Sufi (2015) have advocated for indexation of mortgage principal to a local house price index. Such a design can also reduce the negative effects of housing debt.

5. Summary

Macroprudential policies aimed at crisis prevention and crisis management are primarily about managing the financial-sector and household debt cycles. I posit that minimal innovations to mortgage contracts that create state contingency can deliver considerable social benefits. I reach this conclusion because ex post tools to deal with household mortgage debt are few and face practical implementation challenges. That is, improved mortgage designs are practically the only macroprudential approach to deal with the real economic costs of high levels of household-sector debt. In contrast, the institutional design of central banks as well as the significant impetus underlying post-crisis regulatory reforms provide macroprudential tools to manage financial-sector debt.

References

Admati, A., and M. Hellwig. 2013. *The Bankers' New Clothes: What's Wrong with Banking and What to Do about It.* Princeton, NJ: Princeton University Press.

Agarwal, S., G. Amromin, S. Chomsisengphet, T. Piskorski, A. Seru, and V. W. Yao. 2015. "Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program." Working Paper.

———. 2016. "Policy Intervention in Debt Renegotiation: Evidence from the Home Affordable Modification Program." *Journal of Political Economy* 125 (3): 654–712.

Di Maggio, M., A. Kermani, and R. Ramcharan. 2016. "Monetary Policy Pass-Through: Mortgage Rates, Household Consumption and Voluntary Deleveraging." Working Paper, Harvard Business School.

Duffie, D. 2016. "Financial Regulatory Reform After the Crisis: An Assessment." 2016 ECB Forum on Central Banking, Sintra, Portugal, June 27–29.

Eberly, J., and A. Krishnamurthy. 2014. "Efficient Credit Polices in a Housing Crisis." *Brookings Papers on Economic Activity* (Fall): 73–136.

Fuster, A., and P. S. Willen. 2013. "Payment Size, Negative Equity, and Mortgage Default." NBER Working Paper No. 19345.

Gilchrist, S., and E. Zakrajsek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102 (4): 1692–1720.

Guren, A., A. Krishnamurthy, and T. McQuade. 2016. "Mortgage Design in an Equilibrium Model of the Housing Market." Working Paper, Stanford University.

Jordà, Ò., M. Schularick, and A. Taylor. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking* 45 (s2): 3–28.

Krishnamurthy, A., and T. Muir. 2016. "How Credit Cycles across a Financial Crisis." Working Paper, Stanford University.

Krishnamurthy, A., and A. Vissing-Jorgensen. 2015. "The Impact of Treasury Supply on Financial Sector Lending and Stability." *Journal of Financial Economics* 118 (3): 571–600.

Mian, A., and A. Sufi. 2015. *House of Debt*. 1st ed. University of Chicago Press.

Scholes, M. 2000. "Crisis and Risk Management." *American Economic Review* 90 (2): 17–21.

Leaning Against the Wind When Credit Bites Back*

Karsten R. Gerdrup, Frank Hansen,
Tord Krogh, and Junior Maih
Norges Bank

This paper analyzes the cost-benefit trade-off of leaning against the wind (LAW) in monetary policy. Our starting point is a New Keynesian regime-switching model where the economy can be in a normal state or in a crisis state. The setup enables us to weigh benefits against costs for different systematic LAW policies. We find that the benefits of LAW in terms of a lower frequency of severe financial recessions exceed costs in terms of higher volatility in normal times when the severity of a crisis is endogenous (when "credit bites back"). Our qualitative results are robust to alternative specifications for the probability of a crisis. Our results hinge on the endogeneity of crisis severity. When the severity of a crisis is exogenous, we find that, if anything, it is optimal to lean *with* the wind.

JEL Codes: E12, E52, G01.

1. Introduction

Since the global financial crisis, attention has been devoted to policies that promote financial stability. Empirical literature has found that periods of high credit growth can lead to deeper and more

*This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. The paper was presented at the Riksbank-BIS Workshop "Integrating Financial Stability into Monetary Policy Decision-making—The Modelling Challenges," November 2015; the Annual International Journal of Central Banking Research Conference at the Federal Reserve Bank of San Francisco, November 2016; and at various seminars and workshops in Norges Bank. We are grateful to the participants at the various seminars and workshops, and Lars Svensson, Carl Walsh, Mikael Juselius, Ragna Alstadheim, Andrew Filardo, Phurichai Rungcharoenkitkul, Leif Brubakk, Drago Bergholt, Steinar Holden, Helle Snellingen, and Veronica Harrington for useful comments.

protracted recessions, i.e., “credit bites back” (Jordà, Schularick, and Taylor 2013). Macroprudential policy measures have been introduced in many countries to prevent and mitigate systemic risks related to high growth in credit and asset prices, and can be considered the first line of defense against financial instability since such tools can be both more granular and targeted to the source of risk. However, there is great uncertainty regarding the effectiveness of macroprudential policy tools and the appropriateness of legal frameworks. Monetary policy can still play a powerful role by “leaning against the wind” (LAW), as it is transmitted more broadly to all sectors in the economy and also to shadow banks[1]; see Smets (2014) for a broader discussion.

Financial crises are rare events and have historically occurred every fifteen to twenty years on average; see Taylor (2015). LAW policies can potentially lead to a loss of central bank credibility since such events are hard to predict. Justifying tight monetary policy can be difficult since we cannot observe the possible gains in terms of a lower frequency and severity of crises. If interest rates are systematically kept higher than those implied by price stability, inflation expectations can fall and the inflation-targeting regime could lose credibility. Many authors have furthermore shown that with long-term debt contracts, raising monetary policy rates to reduce the credit-to-GDP level can have perverse effects since GDP growth falls faster and stronger than credit; see Svensson (2014) and Gelain, Lansing, and Natvik (2015). However, gains can be high if leaning against excessive financial stability risks can reduce the probability and severity of crises longer down the road. The recent financial crisis showed that monetary policy, or economic policy in general, has limitations when it comes to cleaning up after a crisis has occurred.

This paper will investigate to what extent monetary policy should react to financial stability risks and deviate from the policy it would have chosen in the absence of such risks. The question of LAW is only relevant insofar as financial stability risk and

[1]Or “it gets in all of the cracks,” as Jeremy Stein put it in a speech at the “Restoring Household Financial Stability after the Great Recession: Why Household Balance Sheets Matter,” a research symposium sponsored by the Federal Reserve Bank of St. Louis, February 7, 2013.

inflation are not perfectly correlated. Negative shocks to inflation may give rise to a trade-off in monetary policy because bringing inflation back to target may have to be achieved at the expense of a buildup of financial imbalances. Small open economies are faced with such trade-offs when international interest rates fall, and the central bank has to reduce domestic interest rates to reduce the degree of currency appreciation and lower imported inflation. The open-economy dimension also gives rise to a trade-off between stabilizing inflation and output when the economy is hit by a shock to aggregate demand (no divine coincidence between stabilizing inflation and output). This means that a central bank cannot clean up after a crisis has erupted by reducing interest rates without paying attention to inflation. For the setup to be relevant for small open economies, we use the estimated model of Justiniano and Preston (2010a) as the core model.

In our framework, LAW is motivated by agents underestimating financial stability risks. When aggregate credit is accumulated in the economy, agents do not incorporate the risk this poses to economic developments. We introduce regime switching into an otherwise standard open-economy New Keynesian model. This procedure enables us to analyze economic developments for an economy that occasionally experiences financial headwinds using a relatively parsimonious model. While the financial system and its interaction with economic developments is extremely complex, we simplify the mechanisms to highlight some important policy trade-offs. We model credit developments as a separate block "outside" the core model. Credit, which will serve as a proxy for financial imbalances in the model, only affects the probability of a crisis and the severity of a crisis and does not affect economic developments in normal times. Parameters controlling the probability of crises and the severity of financial recessions are calibrated based on a sample of OECD countries.

When the economy makes the transition from a normal regime to a crisis regime, aggregate demand is reduced abruptly. The mechanisms behind this aggregate demand shock can be numerous. In Cúrdia and Woodford (2009, 2010) and Woodford (2012) a similar term can be interpreted as credit spreads, the difference in equilibrium yield between long-term bonds issued by risky private borrowers and those issued by the government. Higher credit spreads make financial conditions worse for borrowers, reducing overall welfare. Aggregate demand can also fall as a result of strained balance sheets. When leverage is high, households, non-financial companies, and financial institutions are at higher risk of defaulting. When agents try to de-lever, assets may be sold at fire-sale prices, creating debt-deflation type spirals; see, e.g., Lorenzoni (2008), Bianchi (2011), and Bianchi and Mendoza (2013). Consumers reduce consumption in order to strengthen their balance sheets; see Mian and Sufi (2011) and Dynan (2012). The underlying idea is the same as in Woodford (2012):

> The idea of the positive dependence on leverage is that the more highly levered financial institutions are, the smaller the unexpected decline in asset values required to tip institutions into insolvency—or into a situation where there may be doubts about their solvency—and hence the smaller the exogenous shock required to trigger a crisis. Given some distribution function for the exogenous shocks, the lower the threshold for a shock to trigger a crisis, the larger the probability that a crisis will occur over a given time interval.

This paper is close in spirit to Ajello et al. (2015). They study the intertemporal trade-off between stabilizing current real activity and inflation in normal times and mitigating the possibility of a future financial crisis within a simple New Keynesian model with two states and an endogenously time-varying crisis probability. While they use a two-period setup, we use an infinite time horizon. Using a longer horizon can reduce the benefits of leaning, since credit growth eventually picks up after a monetary policy tightening aimed at mitigating financial stability risk. This point has been highlighted by Svensson (2016) and reflects that real credit, which determines the probability of a crisis, is assumed to return to a specific steady-state level in his application. We calibrate the effect of an interest rate increase on credit growth to reflect SVAR evidence for Norway. Like Ajello et al. (2015), we assume that the crisis probability is a function of credit growth over a five-year period. We follow Alpanda and Ueberfeldt (2016), Pescatori and Laséen (2016), and Svensson (2016) in assuming that a crisis can occur at any point in time. Unlike Ajello et al. (2015) and Alpanda and Ueberfeldt (2016), we assume that the crisis severity is endogenous. This will increase the benefit of leaning in monetary policy. Similar to Ajello et al. (2015), we will assume that private agents underestimate the probability of a crisis. This is in contrast to Alpanda and Ueberfeldt (2016), where agents are rational and perceive aggregate risks correctly, but not their own contribution to that risk.

We contribute to the literature by investigating systematic LAW policies. We find that the benefits of LAW in terms of a lower frequency of severe financial recessions exceed costs in terms of higher volatility in normal times when the severity of crisis is endogenous (when "credit bites back"). The LAW policy can be implemented by responding relatively more to fluctuations in output and/or by responding directly to household credit growth. Compared with a policy that does not take into account that a crisis can happen, a "benign neglect" policy, the LAW policy contributes to a lower loss and reduced tail risk in the economy. The costs are paid in terms of higher inflation volatility. Our qualitative results are robust to alternative calibrations of the probability of crisis. If, however, the severity of crisis is exogenous, or the effect of credit on crisis severity is very small, the optimal response is to lean *with* the wind. This nests the results found in the literature, e.g., Ajello et al. (2015), Alpanda and Ueberfeldt (2016), Pescatori and Laséen (2016), and Svensson (2016), who find no (or very small) net benefits of LAW policies, and typically assume either no or a very small effect from credit to crisis severity.

The rest of the paper is organized as follows. We begin by establishing empirical evidence to inform us about the link between credit and crisis probability and severity in section 2. This is used when we set up a model for analyzing LAW (section 3) and calibrating it (section 4). In section 5 we present some properties of the calibrated model, while the results of our LAW analysis are in section 6. We end with a section on sensitivity (section 7) and conclusions (section 8).

2. Can LAW Policies Make Sense?

We interpret LAW as monetary policy adjustments where the central bank reacts to financial stability risks and deviates from the policy it would have chosen in the absence of such risks. Two central assumptions must be fulfilled for LAW to be able to bring benefits.

Table 1. Estimated Effects of a Monetary Policy Shock on Real Credit in VAR Studies

Paper	Country	Peak Effect of MP Shock (%)
Goodhart and Hofmann (2008)[a]	Panel	Approx. −1.25
Assenmacher-Wesche and Gerlach (2008)[b]	Panel	Approx. −0.8
Musso, Neri, and Stracca (2011)[c]	US	Approx. −3
Musso, Neri, and Stracca (2011)[c]	EA	Approx. −2
Laséen and Strid (2013)[c]	SWE	Approx. −0.8
Robstad (2014)[c]	NOR	Approx. −0.8
Pescatori and Laséen (2016)[c]	CAN	Approx. −0.8

[a]Based on bank loans to the private sector.
[b]Based on total loans to private non-bank residents.
[c]Based on credit to households.

First, monetary policy must be able to affect the relevant financial variables. Second, financial crises cannot be purely exogenous events. The probability of a crisis and/or the severity of a crisis must be linked to financial imbalances. In this section, we evaluate to what extent these necessary conditions are fulfilled, both based on the existing empirical literature and estimates based on a sample of twenty OECD countries. In our setup, we use the five-year growth rate in real household credit as a measure of financial imbalances (similar to, e.g., Ajello et al. 2015).

2.1 The Effect of Monetary Policy on Credit

The empirical literature has established a clear link from monetary policy to credit. Table 1 provides an overview of VAR estimates of the peak effect on the level of real credit following a 1 percentage point increase in the nominal interest rate.

The peak effect of monetary policy on real credit is similar in magnitude across the different studies. This indicates that monetary policy may play a role in affecting credit developments. Papers focusing on (small) open economies (e.g., Laséen and Strid 2013, Robstad 2014, and Pescatori and Laséen 2016) suggest that the peak effect on real household credit following a monetary policy shock that raises the nominal interest rate by 1 percentage point is around 0.8 percent. The effect is estimated to be larger when considering the United States or the euro area as a whole (Musso, Neri, and Stracca 2011). In our paper, we will make use of the estimates in Robstad (2014), which are in the lower end of the estimates.

2.2 The Effect of Credit on the Probability of a Crisis

There is a general consensus in the literature that excessive credit accumulation is one of the main drivers of financial crisis (see, e.g., Reinhart and Rogoff 2008, Schularick and Taylor 2012, Anundsen et al. 2016, and Jordà, Schularick, and Taylor 2016). We estimate the probability of a crisis based on a panel of twenty OECD countries over the period 1975:Q1–2014:Q2.[2] These are the same data used in Anundsen et al. (2016). In contrast to Anundsen et al. (2016), who estimate the probability of being in a pre-crisis period, we reformulate the model and estimate the probability of a crisis start directly (similar to, e.g., Schularick and Taylor 2012). Our dependent variable takes the value of 1 if there was a crisis start in country i at quarter t and zero otherwise.[3] The probability of a crisis start is assumed to depend on the five-year growth in real household credit (L). Using the logit specification, the probability of a crisis start in country i in quarter t is given by

$$p_{i,t} = \frac{\exp(\mu_i + \mu_L L_{i,t})}{1 + \exp(\mu_i + \mu_L L_{i,t})}, \tag{1}$$

where μ_i are country fixed effects and μ_L is the coefficient on the cumulative credit growth. The estimates we get are shown in table 2.

The estimates suggest that the steady-state (annual) probability of a crisis is approximately 3.3 percent. Further, there is a positive effect of credit growth on the probability of a crisis. A plot of the estimated logit function as a function of cumulative real credit growth can be found in figure 8.

[2] Countries included are Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Italy, Japan, Korea, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, the United Kingdom, and the United States.

[3] To avoid post-crisis biases, as discussed in, e.g., Bussiere and Fratzscher (2006), we omit observations for when a given country was in a crisis.

Table 2. Estimated Parameters in the Logit Model

	(1)
Five-Year Real Credit Growth	2.232**
	(1.099)
Constant	−4.792***
	(1.026)
Country Fixed Effects	Yes
Pseudo R-Squared	0.0424
AUROC	0.725
Observations	1,832

Note: Clustered standard errors are reported in parentheses below the point estimates, and the asterisks denote significance levels: * = 10 percent, ** = 5 percent, and *** = 1 percent.

2.3 The Effect of Credit on the Severity of a Crisis

The empirical literature has found that credit accumulation in a boom is important for economic activity during the bust (see, e.g., Jordà, Schularick, and Taylor 2013, 2016 and Hansen and Torstensen 2016). In our framework we are interested in how credit accumulation in the boom affects the path for the output gap during busts. To this end, we use local projection methods (Jordà 2005) to estimate how the five-year growth in real household credit affects the path for the unemployment rate during a financial crisis. We add the unemployment rate to the data used in Anundsen et al. (2016).[4] To keep things simple, we estimate paths only for the unemployment rate during financial recessions (defined as recessions that coincide with a financial crisis within a two-year window).[5] The corresponding effects on the output gap are then calculated based on Okun's law.

[4]Data on unemployment rates were gathered from FRED (Federal Reserve Economic Data). To get as long quarterly series on unemployment rates as possible, the definition of the unemployment rate differs somewhat between countries. For some countries, registered unemployment was used; for others, we have used harmonized rates. Some include the entire population, while some include only the working-age population.

[5]Recession dates are based on Hansen and Torstensen (2016).

Figure 1. Change in Unemployment Rate from the Start of the Crisis (percentage points)

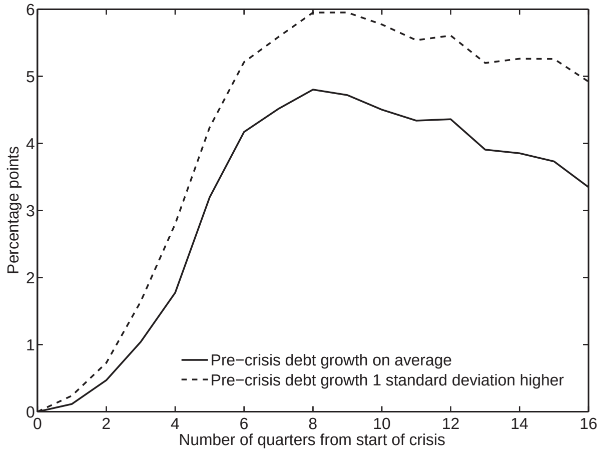

Notes: The solid line is the path for the unemployment rate conditional on the five-year growth in household real credit at the peak of financial recessions. The dashed line illustrates the path when the five-year growth in real household credit is one standard deviation higher at the peak.

Figure 1 and table 3 show the local projection results for the unemployment rate during a financial crisis conditional on pre-crisis cumulative credit growth. In figure 1, the solid black line shows the increase in unemployment during a financial crisis conditional on the average five-year growth in real credit at the peak of the cycle. The dashed line shows the corresponding path for the unemployment rate when the five-year growth in household real credit is one standard deviation higher at the peak.[6] The increase in the unemployment rate during a crisis is both higher and more protracted in this case. The difference between the two paths is also significant; see table 3. Our results are in line with the results in Jordà, Schularick, and Taylor (2013) and suggest that "credit bites back."[7]

[6]The average cumulative growth in real household credit at the peak is 11 percent in the data. The standard deviation is about 17 percent.

[7]The magnitude of this effect is, naturally, uncertain. Estimates may differ depending on the number of countries analyzed, the sample period, and the empirical approach used. For example, Jordà, Schularick, and Taylor (2016) document that the buildup of mortgage credit in the boom has become more important for real economic activity during recessions post–World War II.

Table 3. Local Projection for the Unemployment Rate during Financial Crises

Horizon	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
Constant	0.0	0.3	0.7	1.1	2.5**	3.5***	3.8***	4.1***
	(0.2)	(0.3)	(0.5)	(0.8)	(1.0)	(0.9)	(0.9)	(0.9)
Five-Year Growth in Real Credit	0.7**	1.5**	3.5***	5.9***	6.0***	6.0***	6.2***	6.6***
	(0.3)	(0.5)	(0.9)	(1.5)	(1.6)	(1.5)	(1.5)	(1.6)
Horizon	**(9)**	**(10)**	**(11)**	**(12)**	**(13)**	**(14)**	**(15)**	**(16)**
Constant	3.9***	3.7***	3.6**	3.6*	3.1	3.0	2.8	2.4
	(1.0)	(1.1)	(1.4)	(1.7)	(1.9)	(2.0)	(2.1)	(2.1)
Five-Year Growth in Real Credit	7.1***	7.3***	6.9**	7.2**	7.4**	8.1**	8.8**	9.0**
	(1.8)	(2.0)	(2.4)	(2.9)	(3.2)	(3.4)	(3.6)	(3.6)
Observations	31	31	31	31	31	31	31	31

Notes: Cluster-robust standard errors are in parentheses. Asterisks denote significance: * = 10 percent, ** = 5 percent, and *** = 1 percent.

3. Model for Analyzing LAW

The previous section shows that LAW may bring benefits: monetary policy can affect credit, and credit has an impact on both the probability and the severity of a financial crisis. We therefore introduce a parsimonious framework in order to analyze to what extent monetary policy should conduct LAW policies within a flexible inflation-targeting regime. To achieve this, we add the possibility of large shocks (interpreted as financial stress) to household consumption demand controlled by a Markov process in an otherwise standard New Keynesian open-economy model. In line with the empirical evidence in the previous section, we will make both the probability and the severity of crises endogenous.

3.1 Core Model

As our core model, we will use the small open-economy model in Justiniano and Preston (2010a). The model builds on Galí and Monacelli (2005), Monacelli (2005), and Justiniano and Preston (2010b) and allows for habit formation, indexation of prices, labor market imperfections, and incomplete markets. The reader is referred to Justiniano and Preston (2010a, 2010b) for a detailed description of the model.

3.2 Financial Imbalances

Credit developments play no role in the model of Justiniano and Preston (2010a). Credit is, however, a key variable in our framework, so we add it alongside the core model.

Credit is meant to proxy for financial imbalances, and it will serve two purposes: it will determine endogenously (i) the probability and (ii) the severity of a financial crisis. We will assume that credit is "frictionless" in normal times, which means that there are no direct feedback effects from developments in credit to real economic activity in normal times.

Following Ajello et al. (2015), we let the five-year growth rate in real household credit represent the level of financial imbalances in the model. We denote it L_t, and it is given by

$$L_t = \sum_{s=0}^{19} (\Delta cr_{t-s} - \pi_{t-s}), \tag{2}$$

where Δcr_t is household credit growth and π_t is the inflation rate.

We assume that the growth rate in household credit depends on a vector of endogenous variables (X_t):

$$\Delta cr_t = \omega_X X_t + \epsilon_{\Delta cr,t}, \tag{3}$$

where ω_X is a vector of parameters. $\epsilon_{\Delta cr,t}$ captures shocks to credit.

3.3 Financial Crises

We introduce the possibility of financial crises in the model through Markov switching. A financial crisis is here interpreted as a large, but low-probability, shock to domestic consumption demand. More formally, let $\hat{\epsilon}_{g,t}$ be a shock in the consumption Euler equation (see Justiniano and Preston 2010a for derivations):

$$\begin{aligned} c_t - hc_{t-1} &= E_t(c_{t+1} - hc_t) - \sigma^{-1}(1-h)(i_t - E_t\pi_{t+1}) \\ &\quad + \sigma^{-1}(\hat{\epsilon}_{g,t} - E_t\hat{\epsilon}_{g,t+1}), \end{aligned} \tag{4}$$

where c_t is consumption and i_t is the nominal interest rate. The parameter $\sigma > 0$ is the intertemporal elasticity of substitution and $h \geq 0$ measures the degree of habit in consumption. We will assume that the demand shock consists of two elements: $\hat{\epsilon}_{g,t} = \epsilon_{g,t} - z_t$. $\epsilon_{g,t}$ is a standard autoregressive demand shock, while z_t represents a financial shock:

$$z_t = \rho_z z_{t-1} + \Omega\kappa_t. \tag{5}$$

The parameter Ω is controlled by a Markov process. In normal times, $\Omega = 0$ and $z_t = 0$. In crisis times, $\Omega = 1$ and the crisis impulse κ_t matters for aggregate demand. The crisis impulse is modeled as a function of financial imbalances (L_t):

$$\kappa_t = (1-\Omega)(\gamma + \gamma_L L_t) + \rho_\kappa \Omega\kappa_{t-1}. \tag{6}$$

The parameter γ controls a constant effect on output during a crisis, while the parameter γ_L controls the effect of the initial level of financial imbalances on the severity of a crisis. When the economy is in normal times ($\Omega = 0$), the first term on the right-hand side in equation (6) measures the potential size of the crisis shock, which depends endogenously on developments in financial imbalances (L_t). The probability of switching from normal times to the crisis regime is endogenous and given by $p_{C,t}$, defined as

$$p_{C,t} = \frac{\exp[\mu + \mu_L L_t]}{1 + \exp[\mu + \mu_L L_t]}. \tag{7}$$

The probability of returning to normal times is exogenous and equal to p_N.

3.4 Monetary Policy

The monetary authority has the following loss function:

$$W_0 = E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \lambda_y y_t^2 + \lambda_i (i_t - i_{t-1})^2), \tag{8}$$

where $0 < \beta < 1$ is the household's discount factor. λ_y and λ_i are the weights on the output gap and the change in the nominal interest rate (annualized), respectively.[8] We set $\lambda_y = 2/3$ and $\lambda_i = 1/4$.[9]

We restrict the central bank's interest rate policy to follow optimal simple Taylor-type rules. The exact form of the rules will be spelled out at the relevant stages.

4. Calibration

For calibration, we use the estimates of the parameter values in Justiniano and Preston (2010a).[10] Table 4 shows the calibrated

---

[8]Interest rate changes in the loss function are mainly included to make the dynamics of the interest rate under the optimized policy rules more in line with the interest dynamics under estimated Taylor rules. This term was not included in Gerdrup et al. (2016).

[9]These weights have been used by Norges Bank (see, e.g., Norges Bank 2012).

[10]We have used the median of the posterior in table 2 in their paper.

Table 4. Calibrated Parameters

Parameter	Value	Description
ω_x	0.31	Effect of Output on Credit Growth
ω_r	−0.45	Effect of Real Interest Rate on Credit Growth
$\sigma_{\Delta cr}$	0.014	Standard Deviation of Credit Shock
$\rho_{\Delta cr}$	0.2	Persistence of Credit Shock
ρ_z	0.5	Persistence of Crisis Shock (z_t)
ρ_k	0.8	Persistence in Crisis Impulse(κ_t)
γ	0.55	Exogenous Component in Crisis Severity
γ_L	1.75	Endogenous Component in Crisis Severity
μ	−4.79	Constant Term in the Equation for $p_{C,t}$
μ_L	2.23	Effect of Credit on $p_{C,t}$
p_N	0.125	Probability of Going from Crisis Regime to Normal Times Regime

values for the remaining parameters in the model. This section explains how these parameter values were chosen.

4.1 A "Benign Neglect" Policy Rule

In the baseline calibration, we consider an optimal simple rule based on the assumption that crises never happen. That is, using the loss function specified in equation (8), we search for the optimal simple rule in the original Justiniano and Preston (2010a) model.[11] We label this the "benign neglect" policy rule. It is restricted to take the following form:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)[\theta_\pi \pi_t + \theta_y y_t] + \epsilon_{i,t}, \tag{9}$$

where $\epsilon_{i,t}$ is a monetary policy shock, $0 < \rho_i < 1$ is the degree of interest rate smoothing, and θ_π and θ_y are the response coefficients for inflation and output, respectively.

The optimal parameters in the benign neglect policy rule are $\rho = 0.89$, $\theta_\pi = 6.51$, and $\theta_y = 1.35$. The policy rule features a high degree of interest rate smoothing, and the optimal response to

[11] To simplify, we consider the limiting case when β goes to unity, which transforms the problem to a case where the policymaker minimizes the weighted sum of variances in inflation, output, and interest rate changes.

inflation is relatively higher than the response to output. This rule will serve as a benchmark when we later evaluate policy rules that take the risk of financial crises into account.[12]

4.2 Credit Dynamics

The quarterly rate of household credit growth is assumed to depend on the output gap and the real interest rate. The effect of output and the real interest rate on credit are calibrated in two steps. We first calibrate the effect of output on credit by estimating a simple reduced-form equation for household credit growth.[13] The point estimate of the effect of the output gap on credit growth is reported in table 4. Next, the effect of the real interest rate on credit growth is calibrated to match the response of a monetary policy shock from the structural VAR model in Robstad (2014),[14] given the calibrated effect of output on credit. This ensures that real credit declines about 0.7–0.8 percent in response to a monetary policy shock that raises the nominal interest rate by 1 percentage point at impact.

The credit shock $\epsilon_{\Delta cr,t}$ in equation (3) is assumed to follow an AR(1) process with standard deviation $\sigma_{\Delta cr}$ and persistence $\rho_{\Delta cr}$. We calibrate $\sigma_{\Delta cr}$ and $\rho_{\Delta cr}$ to match (i) the standard deviation in household credit growth and (ii) the correlation between household credit growth and output in Norway.[15]

4.3 The Probability of a Crisis

The probability of a transition from normal times to crisis times ($p_{C,t}$) is assumed to depend on the five-year real growth in household credit (L_t) and is given by (7). We use the estimates documented in section 2.2 to calibrate the parameters μ and μ_L. The probability

[12]The benign neglect rule can be thought of as how policymakers saw the world prior to the global financial crisis from 2007/08.

[13]We regress household credit growth (C2 households) on the output gap (HP-filtered real GDP for mainland Norway using $\lambda = 3,000$). We estimate the model with 2SLS using two lags of the output gap as instruments for the current output gap.

[14]Figure 7 in this paper.

[15]The standard deviation of household credit growth and the correlation with the output gap is empirically (in the model) 1.54 (1.54) and 0.33 (0.38), respectively.

of going from a crisis regime to normal times, p_N, is assumed to be exogenous and calibrated to give an average duration of crisis of two years, requiring $p_N = 0.125$.

4.4 The Severity of a Crisis

In order to calibrate the effect on the output gap in a crisis and how it varies with the pre-crisis credit accumulation, we use the results established in section 2.3. These results indicate that the unemployment rate on average increases by about 5 percentage points during a crisis, while a one standard deviation higher credit accumulation before the crisis increases the unemployment rate by 0.75 percentage points the first two years of the crisis.

We use Okun's law with a parameter of -2 to map the unemployment rate to the output gap. Our results then suggest that the average fall in the output gap during a crisis should be approximately 10 percentage points, while one standard deviation higher credit accumulation before the crisis should make output fall by about 1.5 percentage points more over the first two years of a crisis. With a standard deviation in the five-year growth in real credit of 17 percent, this implies that output falls by $1.5/17 = 0.09$ percentage points more on average during a crisis if the five-year growth in real credit is 1 percentage point higher before the crisis, all else equal. This is in line with the effects of Jordà, Schularick, and Taylor (2013); see the discussion in Svensson (2016, appendix D).

To make the model match the empirical results, we perform local projections on simulated data from the model. The parameters γ and γ_L in equation (6) are then selected to match the results from the empirical projections.[16] The results are shown in figure 2.

[16]Since the output dynamics during a crisis in the model may differ from the implied dynamics illustrated in figure 1, we match the average difference between the output paths over horizon $h = 1, \ldots, H$ in order to capture that the severity of a crisis is both deeper and more protracted when credit growth is high in the boom preceding a crisis. Formally, we let y_h and y_h^L be the implied output paths based on the local projection in figure 1 when credit growth at the peak is on average and one standard deviation higher, respectively, and we let $\hat{y}_h$ and $\hat{y}_h^L$ be the counterparts based on simulations of the model. We choose γ and γ_L to minimize the following loss function: $L(\gamma, \gamma_L) = (\min_h y_h - \min_h \hat{y}_h)^2 + \frac{1}{H}(\sum_{h=1}^{H} (y_h^L - y_h) - \sum_{h=1}^{H} (\hat{y}_h^L - \hat{y}_h))^2$. We use $H = 8$ to be consistent with the assumption of an average duration of crisis of two years.

Figure 2. Output Dynamics during Crises in the Model

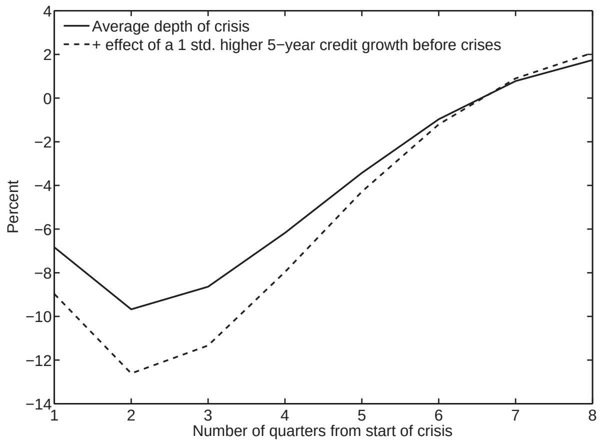

Figure 2 shows that, in the model, the output gap falls by approximately 10 percentage points during a crisis if the five-year real credit growth prior to the crisis is at its average level. Compared with the empirical results in figure 1, the model generates an output path that returns relatively quickly to the pre-crisis level. The average difference between the two paths illustrated in figure 2 is approximately 1.3 percent, somewhat lower than the calibration target of 1.5 percent.

5. Properties of the Calibrated Model

Before we turn to the analysis of policy rules that take the risk of a financial crisis into account, it is useful to consider some properties of the model under the benign neglect policy rule. We ask two questions. First, are there any potential benefits from LAW policies in this economy? Second, how strong is the impact of monetary policy on key variables?

To help answer the first question, table 5 shows statistics from simulations of two versions of the model. First we simulate the benchmark calibrated version, and then a version where crises never happen. The latter simulation can be interpreted as the case where macro stabilization policies have managed to remove crisis risk completely with no side effects to the rest of the economy.

Table 5. Simulations under the Benign Neglect Policy with and without the Possibility of Crises

	Benchmark	If Crises Never Happen
Std. Annual Inflation	1.66	1.59
Std. Output Gap	2.16	1.67
Std. Interest Rate (Ann. %)	3.96	3.17
Std. Real Exchange Rate	8.43	8.19
Std. Credit Growth	1.64	1.58
Loss	100	79.10
Frequency of Crisis (Ann. %)	3.23	0.00

Note: Model standard deviations and the frequency of financial crises are computed by generating 1,000 replications of length 1,000 quarters.

In the benchmark case, the average frequency of crises is about 3.2 percent. The volatility (measured by the standard deviation) of the output gap is reduced by almost one-quarter when crises are removed. The central bank loss is reduced by one-fifth. This implies that crisis risk is an important source of losses for the central bank, and worthy of further analysis. The next section analyzes whether LAW policies can reap any of the benefits from lower crisis risk.

In order to shed light on the second question, we start by plotting the impulse response function (IRF) for key variables to a monetary policy shock; see figure 3.

A monetary policy tightening makes the output gap decline by almost 0.4 percent. Inflation also falls, but it increases in later periods due to the depreciation of the real exchange rate. The five-year real credit growth declines gradually, and the response is much more persistent than for the other variables. The peak impact is at almost –0.8, which was reported as the calibration target in section 4.2.

We also include the IRF for the probability of a crisis. Since this is a non-linear function of credit, the IRF depends on the initial situation of the economy when the monetary policy shock hits.[17]

[17] Here we have simulated the IRF from 50,000 model-generated initial states.

Figure 3. Impulse Response Functions under a Benign Neglect Policy Rule

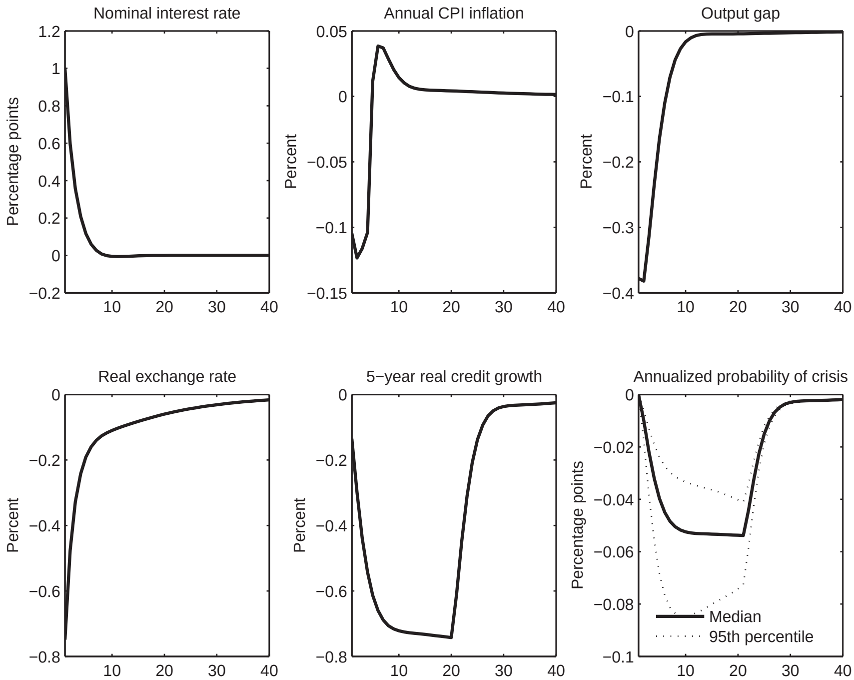

Naturally, the shape of the IRF for the crisis probability resembles the shape of the IRF for the five-year real credit growth. An increase in the policy rate by 1 percentage point leads to a decline in the probability of a crisis of about 0.05 percentage points. This amounts to a decline in p_C from, e.g., 3.2 to 3.15 percent. A larger response can be expected when credit is initially at a higher level. The lower area of the 95th percentile gives a decline of more than 0.08 percentage points.

It is also relevant to check how monetary policy shocks can affect crisis outcomes. Figure 4 shows how a positive monetary shock is expected to affect the output gap if a crisis was to happen sometime in the two-year period after the shock. This also depends on the initial situation of the economy, so it has a distribution across different initial states. The median response tells us that a contractionary

Figure 4. The Effect of a Positive Interest Rate Shock on the Output Gap during a Crisis (percentage points)

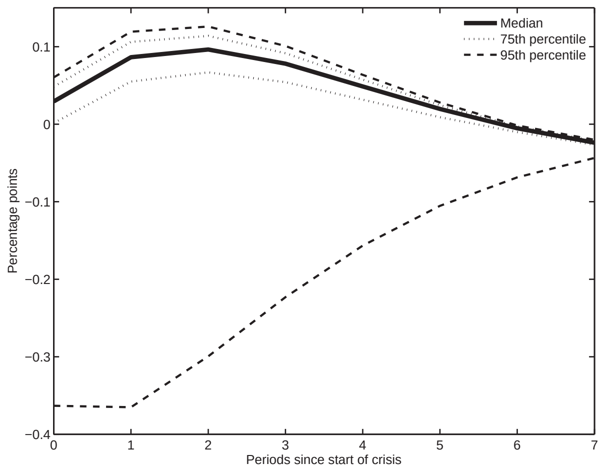

Note: Distribution of crisis outcomes from simulations are based on 50,000 different initial states.

monetary policy shock, through reducing credit, will increase the level of the output gap during crises, i.e., reduce the severity. The largest effect comes in the third period of the crisis, where output is increased by almost 0.1 percentage point. Hence, if the typical fall in output in the third period of the crisis is 10 percentage points, the monetary policy shock lowers it to 9.9 percentage points.

Sometimes the sign shifts, implying that monetary policy makes the output gap even more negative during crises, explaining why the 95th percentile covers area below zero. That happens if a crisis occurs shortly after the monetary policy shock. In such cases the economy is initially weaker due to the monetary policy shock, while the effect of the shock on credit is still quite small.

In summary, monetary policy has the potential to reduce the expected cost of crises, both through a reduction in the probability of a crisis and the potential severity. There will, however, also be a cost associated with such a policy through increased volatility in normal times, especially when there is a trade-off between stabilizing traditional target variables and financial stability.

6. Systematic LAW with Simple Interest Rate Rules

We will now use the model extended with credit and crisis risk to evaluate whether it can be beneficial for the central bank to conduct systematic LAW policies. In particular, we will compare the outcomes under the benign neglect rule and a LAW specification of the following form:

$$i_t = \rho_i i_{t-1} + (1-\rho_i)[\theta_\pi \pi_t + \theta_y y_t + (1-\Omega)\mathbb{1}_{\Delta cr_t > 0}\theta_L(\Delta cr_t - \pi_t)] + \epsilon_{i,t}. \tag{10}$$

θ_L is the response coefficient on real credit growth. We assume an asymmetric response to credit, where $\mathbb{1} = 1$ if real credit growth is positive and zero otherwise. Furthermore, we assume that the central bank does not respond directly to credit growth during crises. The other coefficients have the same interpretation as those in (9).

The motivation for analyzing an asymmetric policy rule is that, in practice, it is natural to think about LAW policies in the context of high credit growth.[18] This means that the interest rate will be kept higher than what is justified by the (medium-term) outlook for inflation and output when credit growth is higher than some threshold. To be pragmatic, we set this threshold to zero (i.e., when real credit growth is above trend).

Table 6 compares the benign neglect policy rule established in section 4 with three different (optimized) policy rules that take the risk of a financial crisis into account.

The first policy rule we consider is labeled *LAW*. In this case, we optimize with respect to all the coefficients in the Taylor rule (10). Introducing endogenous financial crises changes the optimized parameters in several ways. First, the response to output relative to inflation increases and the degree of interest rate smoothing is reduced. Second, the coefficient on credit growth is positive, meaning that monetary policy should react systematically to credit growth.

[18]For example, Norges Bank (2016) states the following: Conditions that imply an increased risk of particularly adverse economic outcomes should be taken into account when setting the key policy rate. This suggests, among other things, that monetary policy should therefore seek to mitigate the *buildup* of financial imbalances.

Table 6. Optimal Parameters in Simple Monetary Policy Rules

Parameter	Benign Neglect	LAW	C-LAW I	C-LAW II
ρ_i	0.89	0.88	0.89	0.86
θ_π	6.51	5.80	6.51	4.60
θ_y	1.35	1.45	1.35	1.24
θ_L	—	0.51	0.64	—

Notes: The optimal coefficients are obtained by minimizing the weighted sum of variances in (annualized) inflation, output, and the change in the nominal interest rate (annualized). The weight on the output gap and the change in the nominal interest rate is $\lambda_y = 2/3$ and $\lambda_i = 1/4$, respectively.

The two remaining policy rules we consider are constrained versions of the LAW policy rule. In the first version, we fix the parameters on the lagged interest rate, inflation, and output in the benign neglect policy rule and reoptimize with respect to credit growth only. We label this policy rule *C-LAW I*. In the second version, we set the coefficient on credit growth to zero and reoptimize with respect inflation, output, and the lagged interest rate (*C-LAW II*). These constrained versions of the LAW policy are meant to illustrate the relative importance of introducing credit growth as an additional element in the Taylor rule and changing the relative response to traditional target variables. While the first constrained version of the LAW policy responds relatively more to credit growth, the second version compensates for the inability to respond directly to credit by increasing the relative response to output. It also features a lower degree of interest rate smoothing.

Table 7 evaluates the different policy rules with regard to the variation in some key variables, the total central bank loss, and the frequency of crises (the unconditional crisis probability). First, comparing the benign neglect policy rule with the different LAW rules, the latter lead to reduced volatility in output but increased costs in terms of higher inflation volatility. The unconstrained LAW policy reduces the total loss by approximately 3.8 percent, and the unconditional probability of a crisis is reduced by 6 basis points. Comparing the two constrained versions of the LAW policy, the

Table 7. Standard Deviations of Endogenous Variables, Loss, and the Frequency of Crisis under Different Policy Rules

	Benign Neglect	LAW	C-LAW I	C-LAW II
Std. Annual Inflation	1.66	1.74	1.69	1.75
Std. Output	2.16	1.83	2.01	1.84
Std. Interest Rate (Ann. %)	3.96	3.94	3.90	3.97
Std. Real Exchange Rate	8.43	8.44	8.41	8.45
Std. Credit Growth	1.64	1.56	1.58	1.58
Loss Relative to Benign Neglect	100.00	96.23	97.62	96.88
Frequency of Crises (Ann. %)	3.23	3.17	3.16	3.21

Note: Model standard deviations and the frequency of financial crisis are computed by generating 1,000 replications of length 1,000 quarters.

results may suggest that taking financial stability considerations into account by changing the relative response to traditional target variables generates a relatively large share of the benefits.

Figure 5 shows the distribution of losses under the different policy rules. While the difference in losses between the alternative LAW policies is small, it is clear that taking crisis risk into account when setting the policy rate reduces the magnitude and frequency of tail-risk events.

In order to illustrate the magnitude of the "degree of leaning" under the alternative LAW policies, we simulate the model under the benign neglect rule, while including the different LAW rules as "cross-checks." Table 8 shows the difference in the nominal interest rate implied by the respective LAW policies and the actual policy rule (benign neglect) for different states of the economy. Considering all states of the economy, the LAW policy implies an interest rate that is approximately 18 basis points higher on average. In periods with elevated financial stability risks (i.e., when $L > 0$), the LAW policy implies a 26 basis points higher interest rate on average. The interest rate is even higher (around 50 basis points) when the real economy is strong at the outset (i.e., $y > 0$). In periods when financial stability risks are elevated but the real economy is weak, the LAW policy implies a somewhat more expansionary policy than the

Figure 5. Distribution of Losses under Different Policy Rules

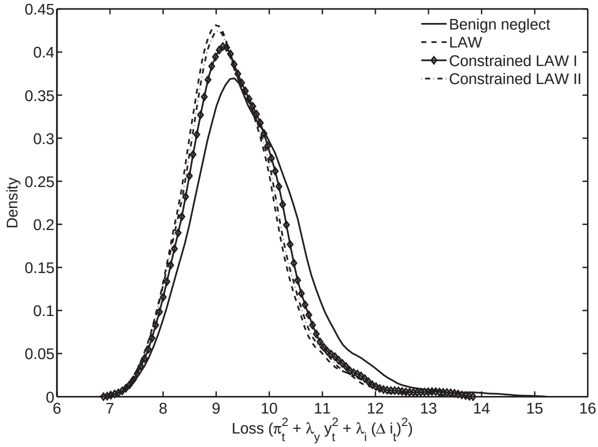

Note: The figure shows the distribution of losses under the alternative policy rules.

Table 8. Difference between the Interest Rate under Alternative LAW Policies and the Benign Neglect Policy Rule for Different States of the Economy (percentage points)

State	Frequency*	LAW	C-LAW I	C-LAW II
All States	1.00	0.18	0.20	0.00
$L > 0$	0.49	0.26	0.25	0.05
$L, y > 0$	0.28	0.48	0.35	0.27
$L > 0, y < 0$	0.21	−0.05	0.12	−0.23
$L < 0$	0.51	0.10	0.16	−0.05
$L, y < 0$	0.29	−0.11	0.08	−0.26
$L < 0, y > 0$	0.21	0.38	0.26	0.23

*The share of time spent in each state.

benign neglect policy. When financial stability risks are relatively low ($L < 0$), the LAW policy implies a 10 basis points higher interest rate. When both financial stability risks are low and the real economy is weak, the interest rate should be 11 basis points lower in the LAW case. The LAW policy rule calls for a higher interest rate (38 basis points) when the real economy is strong ($y > 0$), even though financial stability risks are low.

6.1 A Persistent Decline in Foreign Interest Rates and Financial Stability

The persistent decline in foreign interest rates in recent years has caused a trade-off for many small open economies. In this section we illustrate the dynamics of the economy in normal times and in crisis times when the central bank reacts to the decline in foreign interest rates using the benign neglect rule or the LAW rule. To counter the effects of lower foreign interest rates on the real exchange rate, inflation, and output, a central bank will respond by reducing the domestic policy rate. Lower interest rates might in turn fuel the housing market and increase the accumulation of household credit. This may increase the risk of a financial crisis. In such a scenario, a LAW rule that reduces the interest rate less than what is justified only by the medium-term outlook for inflation and output might deliver better outcomes over time.

Figure 6 shows IRFs of a large and persistent negative shock to the foreign interest rate under the benign neglect rule and the LAW rule. Both policy rules imply a big reduction in the nominal interest rate, but the LAW rule keeps the rate slightly higher. The appreciation of the real exchange rate is greater under the LAW rule and the drop in the inflation rate is larger. On the other hand, the LAW rule stabilizes output more than the benign neglect rule. Together with the higher interest rate, this dampens growth in household credit and reduces both the probability and the potential severity of a financial crisis. Following the negative shock to the foreign interest rate, we impose a crisis after eight quarters (illustrated by the shaded area in figure 6). The LAW policy reduces the severity of a crisis in terms of output. First, it contributes by restraining credit growth prior to the crisis. Second, it responds relatively more to the decline in output and features a lower degree of interest rate smoothing.

The IRFs show what happens when the crisis occurs on a prespecified date. Figure 7 shows the entire distribution for GDP under the two policy rules following the negative shock to the foreign interest rate. The width of the distribution is caused by both (i)

Figure 6. Impulse Responses of a Large and Persistent Negative Shock to Foreign Interest Rates

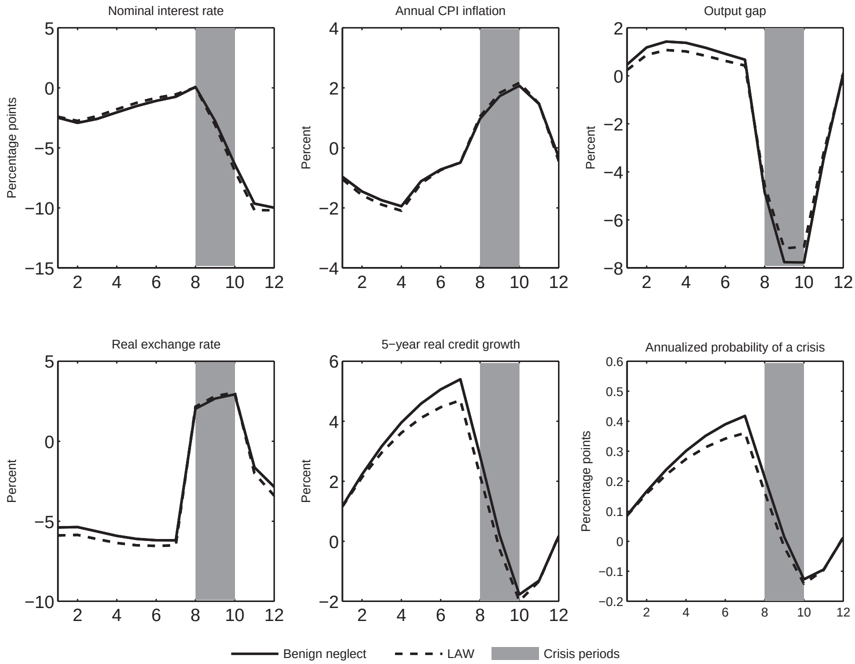

Note: Shaded area indicates a crisis (exogenously imposed after eight quarters).

uncertainty about when the crisis hits and (ii) uncertainty about the severity of the crisis. By following the LAW policy, one reduces the negative tail risk in output by dampening the effect of lower international interest rates on the probability and the potential severity of a crisis.

7. Sensitivity

The parameters governing the probability and the effect of credit accumulation on the severity of a crisis estimated in section 2 are highly uncertain. In this section, we examine how sensitive the optimized response to credit growth in the Taylor rule is to these key parameters.

Figure 7. Distribution of the Output Gap Following a Large, Persistent Negative Shock to Foreign Interest Rate, when Crisis Risk Is the Only Source of Uncertainty after the Interest Rate Shock

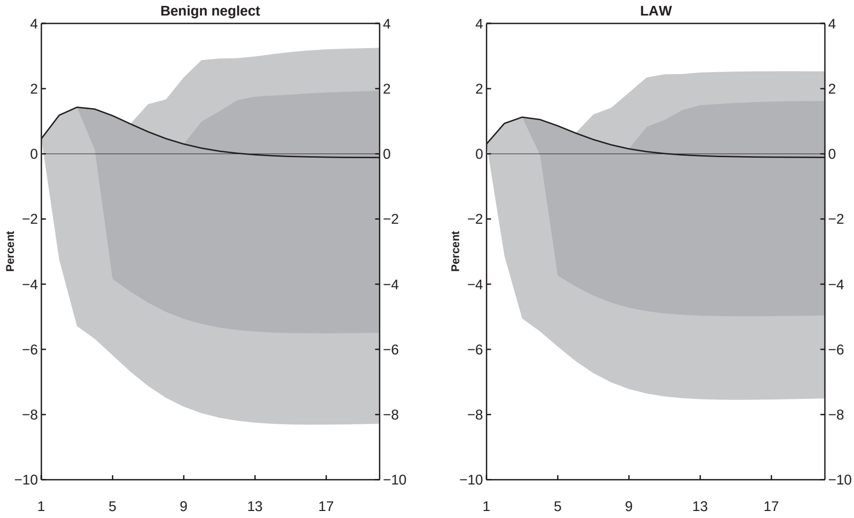

Note: The dark gray area shows the 95th percentile, while the light grey area shows the 99th percentile of the distribution.

7.1 LAW Policies and the Effect of Credit on the Probability of a Crisis

Figure 8 compares the logit function used in this paper with Ajello et al. (2015) and estimates of the effect of mortgage credit on the probability of a crisis for different sample periods reported in Jordà, Schularick, and Taylor (2016).[19] While the steady-state probability used in this paper and the one in Ajello et al. (2015) are similar, the effect of credit on the probability of a crisis is higher in

[19]When plotting the probability functions reported in Jordà, Schularick, and Taylor (2016) (table 4, column 2) we have set the constant term equal to the one used in Ajello et al. (2015). The regressions in Jordà, Schularick, and Taylor (2016) are based on the five-year moving average growth rate in mortgage credit, and not the cumulative growth. This has also been corrected for when plotting the probability functions.

Figure 8. Relationship between Five-Year Real Credit Growth (L) and the Probability of a Crisis Start

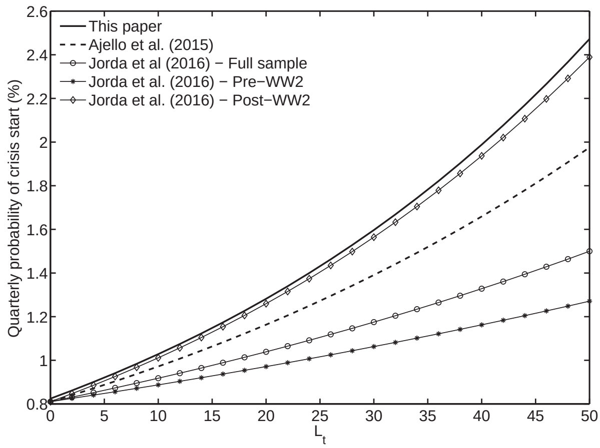

our paper. However, Jordà, Schularick, and Taylor (2016) document that the effect of mortgage credit on the probability of a crisis has increased substantially over time. Estimates used in other papers analyzing LAW policies have typically been based on the long-run historical data in Schularick and Taylor (2012) (e.g., Ajello et al. 2015 and Svensson 2016), which starts in 1870, while we use data from the 1970s. As seen in figure 8, the effect of mortgage credit on the probability used in our paper is close to the effect estimated on the post–World War II sample in Jordà, Schularick, and Taylor (2016).

To illustrate how sensitive the optimized response to credit growth in the Taylor rule is to the underlying assumption about the probability of crisis, table 9 shows optimized parameters for different specifications of the crisis probability. We use the specification reported in Ajello et al. (2015) which features a lower effect of credit on the probability of a crisis. We also examine two cases when the probability of crisis is exogenous: one where the steady-state probability is relatively high (approximately 3.2 percent annually[20]) and one where the probability is relatively low (1 percent).

[20]We have used the steady-state probability reported in Ajello et al. (2015) to facilitate comparison.

Table 9. Optimal Simple Rules for Different Crisis Probabilities

	LAW	Ajello et al. (2015)	Exogenous (3.2% Ann.)	Exogenous (1% Ann.)
ρ_i	0.88	0.87	0.88	0.91
θ_π	5.80	5.17	5.31	6.48
θ_y	1.45	1.30	1.30	1.46
θ_L	0.51	0.28	0.21	−0.03
Loss*	96.23	96.92	98.13	99.77

*Relative to the loss simulated under the benign neglect policy rule (using the same parameters in the logit function for the probability of a crisis).

It is clear that the optimal response to credit growth (and the potential benefit of LAW) depends on the assumed process for the probability of a crisis. First, it depends on the crisis probability level. The coefficient on credit growth is low when the crisis probability is low, both in the case of an endogenous and exogenous crisis probability. In the case of a 1 percent annual exogenous probability of a crisis, the coefficient on credit growth is slightly negative. However, even in this case the relative coefficient on output is somewhat higher than in the benign neglect rule. Second, the steepness of the logit function matters. In Ajello et al. (2015), the crisis probability increases less with accumulated credit growth than in our paper. This is reflected in the coefficient on credit being higher in our benchmark calibration (LAW rule) than that implied by Ajello et al. (2015); see table 9.

While the logit function in our benchmark calibration and the alternative functions shown in this section are linear, or close to linear, introducing more non-linearities may increase the benefits of leaning against the wind to the extent that monetary policy can influence these imbalances. For example, Anundsen et al. (2016) document that the effect of household credit on the likelihood of a crisis increases substantially when it coincides with extreme imbalances (or bubble-like behavior) in the housing market. While such non-linearities (or threshold effects) might be very important, they are difficult to incorporate in the model in a simple way.

Figure 9. The Relationship between the Marginal Effect on GDP during Crisis of a One Standard Deviation Higher Five-Year Growth in Household Real Credit and the Optimal Response to Credit Growth in the Taylor Rule

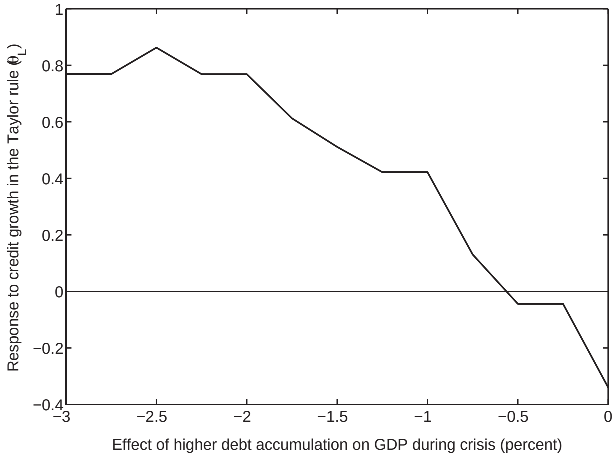

Note: The optimal policy rules are found by searching over an interval of the marginal effect of credit on GDP during a crisis (in steps of –0.25 percent).

7.2 LAW Policies and the Effect of Credit on the Severity of Crisis

The effect of credit accumulation on the decline in GDP during a crisis was calibrated to match the empirical results established in section 2.3. More precisely, our calibration implies that GDP declines by approximately 1.5 percentage points more on average during a crisis if the five-year growth in real household credit before the crisis is one standard deviation higher. In order to illustrate the sensitivity of our results regarding this calibration, figure 9 plots the relationship between the effect of a one standard deviation higher cumulative real credit growth on GDP during crisis (horizontal axis) and the optimal response to credit growth in the Taylor rule.[21] The optimal response to credit is an increasing function of the effect of credit on GDP during a crisis. If we assume that the severity of a crisis is

[21]The other coefficients in the Taylor rule are also reoptimized in this exercise.

exogenous (this implies that the effect of pre-crisis credit growth on GDP during a crisis equals zero), the optimal coefficient on credit growth is negative instead of positive, implying that monetary policy should, if anything, lean *with* the wind. Thus, in a case where the severity of a crisis is exogenous or the effect of credit on GDP during a crisis is sufficiently small, we find no role for monetary policy to respond countercyclically to credit growth.

8. Conclusion

Whether to use monetary policy to curb high growth in credit and asset prices to contain the risk of financial instability, i.e., to "lean against the wind" (LAW), has been the subject of a contentious debate. In this paper we have investigated to what extent monetary policy should actively aim at mitigating the buildup of financial imbalances. LAW is motivated by agents underestimating financial stability risks. We introduce regime switching into an otherwise standard open-economy New Keynesian model to highlight some important policy trade-offs. Credit affects the probability of switching to a crisis and the severity of a crisis, but does not affect economic activity in normal times. Credit is in this sense frictionless in normal times. A transition from a normal regime to a crisis regime involves an abrupt reduction in aggregate demand.

We find that the benefits of LAW in terms of a lower frequency of severe financial recessions exceed costs in terms of higher volatility in normal times when the severity of crisis is endogenous (i.e., "when credit bites back"). The LAW policy can be implemented by responding relatively more to fluctuations in output and by responding directly to household credit growth. Compared with a benign neglect policy, the LAW policy rules contribute to a lower loss and reduced tail risk in the economy. The costs are paid in terms of somewhat higher inflation volatility and interest rate volatility. Our qualitative results are robust to alternative specifications for the probability of a crisis. We also show that the optimal interest rate response to credit growth is higher when the severity of a crisis is sensitive to changes in accumulated credit growth. When the severity of a crisis is exogenous, then it is, if anything, optimal to lean *with* the wind. This nests the results found in the literature, e.g., Ajello et al. (2015), Alpanda and Ueberfeldt (2016), Pescatori and

Laséen (2016), and Svensson (2016), who find no (or very small) net benefits of LAW policies, and typically assume either no or a very small effect from credit to crisis severity. This difference underlines the importance of the assumptions concerning the process which determines the severity of a crisis.

References

Ajello, A., T. Laubach, D. Lopez-Salido, and T. Nakate. 2015. "Financial Stability and Optimal Interest-Rate Policy." Working Paper, Federal Reserve Board (February).

Alpanda, S., and A. Ueberfeldt. 2016. "Should Monetary Policy Lean Against Housing Market Booms?" Working Paper No. 2016/19, Bank of Canada.

Anundsen, A. K., K. Gerdrup, F. Hansen, and K. Kragh-Sørensen. 2016. "Bubbles and Crises: The Role of House Prices and Credit." *Journal of Applied Econometrics* 31 (7): 1291–1311.

Assenmacher-Wesche, K., and S. Gerlach. 2008. "Monetary Policy, Asset Prices and Macroeconomic Conditions: A Panel-VAR Study." Working Paper No. 149, National Bank of Belgium.

Bianchi, J. 2011. "Overborrowing and Systemic Externalities in the Business Cycle." *American Economic Review* 101 (7): 3400–3426.

Bianchi, J., and E. G. Mendoza. 2013. "Optimal Time-Consistent Macroprudential Policy." NBER Working Paper No. 19704.

Bussiere, M., and M. Fratzscher. 2006. "Towards a New Early Warning System of Financial Crises." *Journal of International Money and Finance* 25 (6): 953–73.

Cúrdia, V., and M. Woodford. 2009. "Credit Frictions and Optimal Monetary Policy." BIS Working Paper No. 278.

———. 2010. "Credit Spreads and Monetary Policy." *Journal of Money, Credit and Banking* 42 (s1): 3–35.

Dynan, K. 2012. "Is a Household Debt Overhang Holding Back Consumption?" *Brookings Papers on Economic Activity* 44 (1, Spring): 299–362.

Galí, J., and T. Monacelli. 2005. "Monetary Policy and Exchange Rate Volatility in a Small Open Economy." *Review of Economic Studies* 72 (3): 707–34.

Gelain, P., K. J. Lansing, and G. J. Natvik. 2015. "Leaning Against the Credit Cycle." Working Paper No. 4/2015, Norges Bank (February).

Gerdrup, K., F. Hansen, T. Krogh, and J. Maih. 2016. "Leaning Against the Wind When Credit Bites Back." Working Paper No. 9/2016, Norges Bank.

Goodhart, C., and B. Hofmann. 2008. "House Prices, Money, Credit, and the Macroeconomy." *Oxford Review of Economic Policy* 24 (1): 180–205.

Hansen, F., and K. N. Torstensen. 2016. "Does High Debt Growth in Upturns Lead to a More Pronounced Fall in Consumption in Downturns?" Economic Commentaries No. 8/2016, Norges Bank.

Jordà, O. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–82.

Jordà, O., M. Schularick, and A. M. Taylor. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking* 45 (s2): 3–28.

———. 2016. "The Great Mortgaging: Housing Finance, Crises and Business Cycles." *Economic Policy* 31 (85): 107–52.

Justiniano, A., and B. Preston. 2010a. "Can Structural Small Open-Economy Models Account for the Influence of Foreign Disturbances?" *Journal of International Economics* 81 (1): 61–74.

———. 2010b. "Monetary Policy and Uncertainty in an Empirical Small Open-Economy Model." *Journal of Applied Econometrics* 25 (1): 93–128.

Laséen, S., and I. Strid. 2013. "Debt Dynamics and Monetary Policy: A Note." Working Paper No. 283, Sveriges Riksbank.

Lorenzoni, G. 2008. "Inefficient Credit Booms." *Review of Economic Studies* 75 (3): 809–33.

Mian, A., and A. Sufi. 2011. "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis." *American Economic Review* 101 (5): 2132–56.

Monacelli, T. 2005. "Monetary Policy in a Low Pass-Through Environment." *Journal of Money, Credit and Banking* 37 (6): 1047–66.

Musso, A., S. Neri, and L. Stracca. 2011. "Housing, Consumption and Monetary Policy: How Different Are the US and the Euro Area?" *Journal of Banking and Finance* 35 (11): 3019–41.

Norges Bank. 2012. "Monetary Policy Report 3/12." Technical Report No. 3.

———. 2016. "Monetary Policy Report with Financial Stability Assessment 3/16." Technical Report No. 3.

Pescatori, A., and S. Laséen. 2016. "Financial Stability and Interest-Rate Policy: A Quantitative Assessment of Costs and Benefits." IMF Working Paper No. 16/73 (March).

Reinhart, C. M., and K. S. Rogoff. 2008. "Is the 2007 US Sub-prime Financial Crisis So Different? An International Historical Comparison." *American Economic Review* 98 (2): 339–44.

Robstad, Ø. 2014. "House Prices, Credit and the Effect of Monetary Policy in Norway: Evidence from Structural VAR Models." Working Paper No. 5/2014, Norges Bank.

Schularick, M., and A. M. Taylor. 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008." *American Economic Review* 102 (2): 1029–61.

Smets, F. 2014. "Financial Stability and Monetary Policy: How Closely Interlinked?" *International Journal of Central Banking* 10 (2): 263–300.

Svensson, L. E. O. 2014. "Inflation Targeting and 'Leaning Against the Wind'." *International Journal of Central Banking* 10 (2): 103–14.

———. 2016. "Cost-benefit Analysis of Leaning Against the Wind: Are Costs Larger Also with Less Effective Macroprudential Policy?" NBER Working Paper No. 21902 (January).

Taylor, A. M. 2015. "Credit, Financial Stability, and the Macroeconomy." NBER Working Paper No. 21039 (March).

Woodford, M. 2012. "Inflation Targeting and Financial Stability." NBER Working Paper No. 17967 (April).

Discussion of "Leaning Against the Wind When Credit Bites Back"*

Carl E. Walsh
University of California, Santa Cruz

1. Introduction

The 2008–09 global financial crisis and the ensuing Great Recession forced monetary policymakers to rethink the role of financial market imbalances and asset price developments in the design of monetary policy. It has led economists to ask whether, in the buildup to the crisis, the Federal Reserve and other central banks should have responded to the rapid increase in housing prices and in debt by raising interest rates more than would normally be justified based on the behavior of inflation and real economic activity. That ex post reexamination of past policy decisions has been accompanied by a debate over whether, in the post-crisis era, policy should actively respond to credit growth by "leaning against the wind."

Gerdrup et al. (this issue) directly address the desirability of leaning against the wind, or LAW, policies. Following Svensson (2016), such policies mean "conducting, for financial stability purposes, a tighter monetary policy (a higher interest rate) than justified by standard flexible inflation targeting when the possibility of a financial crisis is disregarded." Defining LAW policies carefully is important, as a flexible inflation-targeting central bank will react to financial market developments to the extent they are useful for forecasting future developments in inflation or real economic activity.[1] The relevance of financial conditions for monetary policy was actually recognized in the Federal Reserve Reform Act of 1977, which states that "the Board of Governors of the Federal Reserve System and the Federal Open Market Committee shall maintain *long run*

*Prepared for the November 2016 International Journal of Central Banking Conference held at the Federal Reserve Bank of San Francisco. Author contact: walshc@ucsc.edu.

[1]See the discussion in Walsh (2009).

growth of the monetary and credit aggregates commensurate with the economy's long run potential to increase production, so as to promote effectively the goals of maximum employment, stable prices and moderate long-term interest rates." (emphasis added)

LAW policies are controversial. Several authors (e.g., see Gambacorta and Signoretti 2014, Filardo and Rungcharoenkitkul 2016) have argued that the benefits of LAW policies outweigh any costs.[2] This view has influenced policy in Norway, for example, where Norges Bank Governor Øystein Olsen (2015) has stated that "a robust monetary policy should therefore take into account the risk of a build-up in financial imbalances." Others (for example, Ajello et al. 2016) find that accounting for financial crisis concerns causes optimal policy to deviate very little from the case in which such concerns are ignored, and that optimal policy may, in some cases, actually lean *with* the wind.[3] The most prominent opponent of LAW policies has been Lars Svensson, who has argued that the costs of such policies far outweigh any potential benefit; see Svensson (2016, 2017).[4]

The broad outlines of the cost-benefit analysis of leaning against the wind are straightforward. Raising the policy rate in response to rapid credit growth may reduce such growth, reduce the probability of a crisis, and lessen the consequences of a crisis should one actually occur. These are the benefits. The costs take the form of the reduced economic activity and shortfall of inflation relative to target that are a result of the tighter policy. Given that there are both costs and benefits, the issues at debate are primarily empirical, and the challenge is to quantify accurately these costs and benefits.

Gerdrup et al. (this issue) offer a new contribution to this debate over costs and benefits. Specifically, they extend the two-period model of Ajello et al. (2016) to the infinite-horizon, small open-economy case. Importantly, they endogenize both the probability of

[2] According to Filardo and Rungcharoenkitkul (2016), "the shortcomings of macroprudential tools have left open an important role for monetary policy to lean against the wind."

[3] Ajello et al. (2016) find that the optimal extent of leaning against the wind is somewhat larger in the face of uncertainty and if the policymaker desires robustness.

[4] Other contributors to the debate include Agur and Demertzis (2013) and Gelain, Lansing, and Natvik (2015).

a financial crisis and the severity of the crisis. These aspects of the model are critical for any evaluation of LAW policies, as I discuss below. With their model in hand, they investigate the optimal coefficient on credit growth in a Taylor-type instrument rule for monetary policy. They also study the distribution of real output and inflation under the optimal LAW instrument rule.

In my comments on their paper, I begin by reviewing the costs and benefits of LAW policies. Doing so will highlight the likely importance of some channels over others. It will also identify the channels the authors incorporate into their analysis and a couple that they, like the rest of this literature, do not. I discuss the approach they take to evaluating the consequences of LAW policies and summarize their primary conclusions. I then offer some comments on their calibration approach, the specification of policy, and potential limitations of the general approach.

2. Costs and Benefits

To organize a discussion of the costs and benefits of LAW policies, suppose the economy can be in one of two states. The first, denoted by subscript NC, is the non-crisis state; the second, denoted by subscript C, is the crisis state. Let L_i denote the present discounted value of losses if the economy is in state $i = NC, C$, and let l_i be the current-period loss in state i. To complete the notation, let β be the discount rate, p_C the probability of moving from a non-crisis to a crisis state, and p_N the probability of moving from the crisis to a non-crisis state.

Then L_{NC} and L_C are defined by

$$L_{NC} = l_{NC} + \beta \left[p_C L_C + (1 - p_C) L_{NC} \right] \tag{1}$$

and

$$L_C = l_C + \beta \left[p_N L_{NC} + (1 - p_N) L_C \right]. \tag{2}$$

Before proceeding, it is important to recognize that this simple specification ignores much. For example, it ignores the fact that L_{NC} (L_C) may depend on the previous state. Even so, (1) and (2) can help organize a discussion of the costs and benefits of LAW policies.

Consider the effects of an increase in the nominal interest rate i_{NC} in the non-crisis state and its effects on current and future losses. By solving (1) and (2) jointly, the effect on L_{NC} of a change in i_{NC} is given by

$$\frac{dL_{NC}}{di_{NC}} = \alpha_1 \frac{dl_{NC}}{di_{NC}} + \alpha_2 \frac{dl_C}{di_{NC}} + (L_C - L_{NC}) \left(\alpha_3 \frac{dp_C}{di_{NC}} - \alpha_4 \frac{dp_N}{di_{NC}} \right), \tag{3}$$

where the constants α_j are positive and functions of β, p_C, and p_N. The terms on the right of (3) summarize the channels through which an interest rate increase motivated by financial stability concerns causes the present value of losses to fall or rise, i.e., whether dL_{NC}/di_{NC} is negative or positive. If it is negative, the benefits outweigh the costs and leaning against the wind is desirable; it if is positive, costs outweigh benefits.

The first term, $\alpha_1 dl_{NC}/di_{NC}$, is the change in the loss in the non-crisis state. This is generally assumed to be positive—by raising the policy rate in the non-crisis state, economy activity will slow and inflation will decline, worsening outcomes in the non-crisis state. The second term, $\alpha_2 dl_C/di_{NC}$, measures the effect on losses conditional on being in a crisis. This effect could be positive (a cost) or negative (a benefit). Svensson (2016) has emphasized that the costs of a crisis are larger if the economy enters it in a weaker condition such as would be the case if the policy rate had been increased prior to a crisis. In contrast, if leaning against the wind before the crisis helps limit the growth of financial imbalances, it might also limit the severity of a crisis, should a crisis occur. This endogenous severity channel is one of the potential benefits of LAW policies that Gerdrup et al. stress.

The final two terms capture the effects of policy on the transition probabilities. If leaning against the wind reduces the probability of a crisis, $dp_c/di_{NC} < 0$, and, because $L_C > L_{NC}$, this reduces the present value of losses (a benefit). Similarly, if leaning against the wind increases the probability of exiting from a crisis and leads on average to shorter crises, than $dp_N/di_{NC} > 0$, which adds to the benefits of LAW policies.

Suppose $\beta = 0.96$ (annual frequency), $p_C = 0.032$, and $p_N = 0.5$.[5] The values of the α_i parameters can be evaluated and (3) becomes

$$\frac{dL_{NC}}{di_{NC}} \approx 23.61\frac{dl_{NC}}{di_{NC}} + 1.39\frac{dl_C}{di_{NC}} + (L_C - L_{NC})\left(22.66\frac{dp_C}{di_{NC}} - 1.34\frac{dp_N}{di_{NC}}\right). \tag{4}$$

While representing the outcome of a simple exercise, (4) offers some suggestive implications. First, the coefficient $dl_{NC}/d_{i_{NC}}$, which is the impact of leaning against the wind on non-crisis welfare, is large. With crises infrequent, the loss in the non-crisis period of tighter monetary policy is likely to loom large in the cost-benefit calculation. Second, and again because crises are infrequent, the potential gain in a reduced crisis loss from LAW policies, measured by dl_C/di_{NC}, receives a much smaller weight than does dl_{NC}/di_{NC}. Third, the impact of LAW policies on the probability of a crisis is more important than its impact on the probability of exiting a crisis. However, with p_C much smaller than p_N, dp_C/di_{NC} and dp_N/di_{NC} are unlikely to be of similar magnitude. Letting e_j denote the elasticity of p_j with respect to i_{NC}, for $j = C, N$, and using the calibrated values $p_C = 0.032$ and $p_N = 0.5$,

$$\begin{aligned}\left(22.66\frac{dp_C}{di_{NC}} - 1.34\frac{dp_N}{di_{NC}}\right) &= \left(\frac{1}{i_{NC}}\right)(22.66e_C p_C - 1.34e_N p_N) \\ &= \left(\frac{1}{i_{NC}}\right)(0.72e_C - 0.67e_N).\end{aligned} \tag{5}$$

Thus, a policy that reduces the probability of a crisis by 10 percent (from 3.2 percent per year to 2.9 percent) would have similar benefit to a policy of leaning against the wind that, by leading to a more shallow crisis, increased the exit probability by 10 percent (from 50 percent per year to 55 percent). Finally, the potential gains from a LAW policy are increasing in the size of the loss during a crisis event. If crises are very costly—i.e., if $L_C - L_{NC}$ is large—then a

[5] These values are consistent with the calibration of Gerdrup et al. (this issue) and the evidence in Jordà, Schularick, and Taylor (2013, 2017).

large effect on p_C might dominate the non-crisis loss represented by dl_{NC}/di_{NC}.

The model employed by Gerdrup et al. (this issue) incorporates three of the four channels included in (3); the only one they do not endogenize is the crisis exit probability. By ignoring the exit probability channel, they follow the standard in this literature. However, the simple calculation leading to (5) suggests that policies that increase the probability of exiting from the crisis state may be as valuable as policies that decrease the probability of entering a crisis by the same percent. Jordà, Schularick, and Taylor (2013) note that "financial crisis recessions are costlier, and more credit-intensive expansions tend to be followed by deeper recessions . . . *and slower recoveries.*" (emphasis added). Thus, a benefit of LAW policies may be the impact they have on shortening the recession associated with a crisis.

Setting aside any effects on p_N, (3) becomes

$$\frac{dL_{NC}}{di_{NC}} = \alpha_1 \frac{dl_{NC}}{di_{NC}} + \alpha_2 \frac{dl_C}{di_{NC}} + (L_C - L_{NC})\,\alpha_3 \frac{dp_C}{di_{NC}}.$$

The first term on the right is considered the cost of LAW policies and the third term is a benefit; the second term, dl_C/di_{NC}, is more controversial, as it may be positive (a cost) or negative (a benefit). Svensson has emphasized that if the economy enters a crisis in a weaker state, losses during the crisis will be larger. Gerdrup et al. (this issue) allow for this channel in their model, as the loss during the crisis depends on the economy's state when a crisis occurs. However, they also allow for LAW policies to contribute to a milder crisis by endogenizing the severity of the crisis. In fact, this is one of the major differences in their analysis from that of Ajello et al. (2016). The authors find that this endogenous severity channel is critically important for generating a role for LAW. In fact, consistent with Ajello et al. (2016) and Svensson (2016), the authors find that if crisis severity is exogenous, the optimal policy rule actually leans *with* the wind.

One further point is relevant before getting to the specifics of the authors' contributions. In addition to ignoring effects on exit probabilities, the analysis of Ajello et al. (2016) and Gerdrup et al. (this issue) assumes that LAW policies increase loss in the non-crisis state; that is, they assume dl_{NC}/di_{NC} is positive. This accords with the evidence in Jordà, Schularick, and Taylor (2013), who find, during the post–World War II era, that "excess credit growth appears to be associated with longer periods of economic growth . . . and expansions last almost 5 years longer in periods of high excess credit growth." Any LAW policy that slows credit growth must shorten the pre-crisis period of economic expansion. However, if the credit boom reflects distortionary growth, a longer expansion may simply magnify the distortions associated with the boom. Such an expansion need not translate into an increase in welfare. For example, the rapid growth in housing investment in the United States in the period leading up to the global financial crisis left a legacy of abandoned, zombie housing developments in its wake. If this reflects a misallocation of capital, then the resulting efficiency losses might imply that the extended boom was not consistent with the Federal Reserve Act's call for credit growth "commensurate with the economy's long run potential to increase production."

3. The Model

Now let me turn to the specifics of the authors' model. As already mentioned, it extends the two-period framework of Ajello et al. (2016) to an infinite-horizon model of a small open economy. To highlight the basic structure, though, I will ignore the open-economy aspects.[6] In this case, a stripped-down version of their model consists of an expectational Euler equation given by

$$x_t = \mathrm{E}_t x_{t+1} - \sigma^{-1} \left(i_t - \mathrm{E}_t \pi_{t+1} - \varepsilon_t \right), \tag{6}$$

where[7]

$$\varepsilon_t \equiv \left(\epsilon_{g,t} - \mathrm{E}_t \epsilon_{g,t+1} \right) - \left(z_t - \mathrm{E}_t z_{t+1} \right),$$

and a New Keynesian Phillips curve given by

$$\pi_t = \beta \mathrm{E}_t \pi_{t+1} + \kappa x_t. \tag{7}$$

[6]Their model also includes habit persistence and partial indexation of prices, and distinguishes between domestic and CPI inflation, output and domestic consumption, imperfect pass-through, and deviations from uncovered interest parity condition. These additions lead to a model better able to match data but are not central to the issues I want to raise.

[7]My notation differs slightly from that of the authors.

A crisis is represented by a positive realization of z_t, which constitutes a negative aggregate demand shock. Specifically, the financial shock z_t is governed by

$$z_t = \rho_{z,t} z_{t-1} + \Omega \kappa_t$$
$$\kappa_t = (1 - \Omega)(\gamma + \gamma_L L_t) + \rho_\kappa \Omega \kappa_{t-1},$$

where $\Omega = 0$ in the non-crisis regime, $\Omega = 1$ in the crisis regime, and L_t is the cumulative five-year growth of real, household credit.

Consider what happens if the economy begins in a steady-state, non-crisis state with $\Omega = 0$ and $z_t = 0$. If a crisis occurs, the size of the negative shock to z_t equals $\gamma + \gamma_L L_t$ and is increasing in the extent to which credit has grown during the previous five years. Cumulative real credit growth is, in turn, assumed to depend on the output gap, inflation, and the real rate of interest (and therefore on monetary policy). The dependence of the severity of the crisis, as measured by κ_t, on pre-crisis credit growth and monetary policy is one channel through which a LAW policy can affect outcomes during a crisis. The second channel Gerdrup et al. (this issue) incorporate is to allow credit growth to affect the probability that a crisis occurs. In normal times, there is no feedback from credit growth to either the real economy or inflation unless monetary policy reacts to credit growth. What credit growth does affect is the probability that a crisis occurs and the severity of a crisis when it does occur.

3.1 Policy

The central bank is assumed to employ an asymmetric version of a Taylor rule:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)\left[\theta_\pi \pi_t + \theta_x x_t + (1 - \Omega)\,\theta_L \left(\Delta cr_t - \pi_t\right)_{\Delta cr_t \geq 0}\right], \tag{8}$$

where Δcr_t is nominal credit growth. In this formulation, the central bank only responds to growth in real credit in non-crisis periods and then only when credit growth is positive.[8] In the context of

[8]Gerdrup et al. (this issue) also introduce a monetary shock into (8), but that is not relevant for discussing LAW policies and so I ignore it.

(8), a policy of leaning against the wind is interpreted to mean that $\theta_L > 0$.

To determine whether a LAW policy is desirable, the authors postulate that loss depends on volatility in inflation, the output gap, and the change in the nominal rate of interest. Thus, the coefficients in (8) are chosen to minimize

$$W_0 = \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\pi_t^2 + \lambda_x x_t^2 + \lambda_i \left(\Delta i_t \right)^2 \right], \tag{9}$$

with $\lambda_x = 2/3$, $\lambda_i = 1/4$ when inflation is expressed at an annual rate. Importantly, credit does not appear in the loss function. If it is optimal to respond to credit growth (i.e., if $\theta_L > 0$), this will not be due to any inherent desire to stabilize credit growth but will instead mean that responding to credit growth helps better achieve objectives defined in terms of inflation, the output gap, and interest rate changes.

The results on the optimized coefficients in (8) will depend on the relative weights put on the terms in the loss function. By way of comparison, therefore, it is worth noting that in their similar exercise, Ajello et al. (2016) assume $\lambda_x = 1/16$, $\lambda_i = 0$. Thus, relative to Ajello et al., Gerdrup et al. put more weight on output gap stabilization and on reducing volatility in nominal interest rate changes.

4. Calibration and Results

The authors draw on the existing literature and their own estimation to calibrate model parameters. Some parameters are standard, but the critical aspects of the exercise are these related to the sensitivity of credit growth to monetary policy, crisis probability to credit growth, and crisis severity to credit growth. The first of these (monetary policy effects on credit growth) is estimated using data from Norway; the second and third are based on data from twenty-two OECD countries. I will come back to this issue below.

The basic policy experiment assumes myopic private agents in the sense that the central bank knows the true crisis process, but private agents perceive the probability of a crisis to be zero. In this environment, the optimal instrument rule is found under three alternative specifications: (1) the central bank is also myopic in ignoring the possibility of a crisis and does not respond to credit growth (labeled the benign neglect case); (2) the central bank takes into account the possibility of a crisis but does not directly respond to credit growth (i.e., it maintains $\theta_L = 0$); and (3) the LAW case in which θ_L in (8) can differ from zero. The optimal policy coefficients for these three cases are reported in table 3 of the paper. Taking into account the possibility of a crisis but without responding directly to credit (moving from case 1 to case 2) has two effects on the optimized instrument rule. First, the responses to both inflation and output are muted. For example, the coefficient on inflation falls from 6.51 to 4.42 while that on output falls slightly from 1.35 to 1.11. Second, the response to output relative to inflation increases. Table 4 of the paper shows that the benign neglect policy produces more stable inflation and interest rates but more volatile output relative to the constrained LAW policy.

The key results are those for case 3. The optimal responses to inflation (5.63) and output (1.25) fall between the benign neglect and constrained LAW cases, but, more importantly, the optimal value of θ_L is positive (0.61). Policy leans against the wind. Comparing outcomes in table 4 under the LAW policy and the constrained LAW policy, credit crises are marginally less likely to occur (the annual probability of a crisis falls from 3.28 percent to 3.17 percent), and as shown in figure 5 of the paper, the distribution of losses given by (9) is shifted to the left and the likelihood of large losses is reduced under the LAW policy. When credit growth is positive, the nominal interest averages 28 basis points higher under the LAW policy.

5. Comments

The authors' results provide a very useful contribution to the LAW debate. My comments will be directed to three issues: the parameter values employed for the calibration exercises, the specification of policy, and the potential limitations of the general approach.

5.1 Consistency of Parameters

Using an empirically grounded model to investigate the desirability of LAW policies is commendable, and it is an important contribution.

Employing a small open-economy framework is also a nice contribution. To assign values of the model's parameters, however, the authors draw on a number of sources. The basic parameters of the open-economy model are drawn from Justiniano and Preston (2010), who estimate a DSGE model using data from Australia, Canada, and New Zealand. The equation describing the evolution of credit and the impact of monetary policy on credit growth is estimated using data from Norway. The relationship between credit growth and both the probability of a financial crisis and its severity is estimated by extending the work of Anundsen et al. (2016) to include data from twenty-two OECD countries over the period 1975:Q1–2014:Q2. None of these choices is unreasonable, but it does open up the question of exactly how realistic the resulting mix is as a representation of a generic open economy. Australia, Canada, New Zealand, and Norway are all exporters of non-manufactured goods, primarily commodities and agricultural products; for other small open economies, the calibration based on Justiniano and Preston (2010) may be less realistic.

The empirical model of credit growth for Norway is significantly different from that employed by Ajello et al. (2016) based on U.S. data, and the differences are likely to be important. For example, Ajello et al. link nominal credit growth in the United States to the output gap, inflation, and the nominal interest rate and obtain

$$\Delta cr_t = 0.18x_t - 0.26i_t + 1.43\pi_t. \tag{10}$$

Based on Norwegian data, Gerdrup et al. obtain

$$\Delta cr_t = 0.31x_t - 0.79\left(i_t - \mathrm{E}_t\pi_{t+1}\right). \tag{11}$$

Conditional on the other variables, a rise in the nominal interest rate has a much larger contractionary effect on credit growth in a model based on (11) then if the Ajello equation were used. This matters for the analysis; all else equal, the marginal benefit of a LAW policy will be larger in the calibrated model based on (11) than if it were based on (10), as a smaller rise in the policy rate is needed to achieve a given reduction in credit growth.

For a small open-economy model, the estimates based on Norway are perhaps more relevant then those based on the United States. But the critical linkages between credit growth, the probability of a crisis, and the severity of a crisis come from a larger sample of OECD countries that includes, for example, the United States. In this case, there is the question of whether the estimated relationships are the right ones to employ for the small open-economy setting the authors have in mind.

Fortunately, the authors do conduct some robustness checks, and the general qualitative conclusions appear robust. But it would be nice—and this is along the lines of a suggestion for future research—to have a model calibrated consistently to represent a specific country (such as Norway). This could be done for most of the parameters using existing models that have been estimated for various small open economies. The hardest parameters to pin down for an individual country would be those related to the impact of credit growth on the probability and severity of a crisis.

A second issue also related to the calibration is the question of the economic significance of the results. The authors find clear support for LAW policies, and their findings emphasize the importance of incorporating feedback from policy to credit growth to the probability and severity of a crisis. But the quantitative magnitude of the effects they find seems small. For example, the policy rate averages 28 basis points higher under the optimal LAW policy. A natural question is whether this magnitude is statistically significant once parameter estimation error is factored in. How confident can we be that the hypothesis that central banks should not lean against the wind can be rejected?

5.2 The Policy Framework

The core of the authors' contribution is their investigation of the optimal value of the policy parameter θ_L in the instrument rule (8). The authors assume policy is implemented via a rule whose specific form is given by (8). To determine how policy should respond to credit growth, they need a metric by which to rank outcomes for different values of the response coefficients in (8). This ranking is done by using the objective function given in (9). But given the objective function (9), one could dispense with (8) and examine optimal policy, i.e., the policy that minimizes (9) without the additional constraint that the resulting description of policy coincide with (8).

Of course, there is a huge literature examining optimized instrument rules. But given that the Norges Bank has been at the forefront in deriving policy from a specification of the objective function rather then by assuming a rule, it is perhaps surprising that the authors took the approach they did. However, there are some advantages to focusing on simple rules. The LAW debate is about whether central banks should respond to credit growth, and focusing on the response coefficients in (8) allows one to have a direct answer to the question simply by determining whether outcomes (valued in terms of (9)) are improved for positive values θ_L.

However, one would like to understand better the extent to which a policy recommendation is inherent to the economic environment in which the probability and severity of a financial crisis are endogenous. A danger with simple rules is that credit growth may appear in an optimized instrument rule because the chosen rule is suboptimal and the presence of credit growth allows the rule to approximate more closely an optimal policy.

To illustrate this issue, consider a simplified version of the model in which the interest rate smoothing objective in (9) is ignored. In this case, the policy problem is

$$\min \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \lambda_x x_t^2\right)$$

subject to

$$x_t = \mathrm{E}_t x_{t+1} - \sigma^{-1}\left(i_t - \mathrm{E}_t \pi_{t+1} - \varepsilon_t\right),$$
$$\pi_t = \beta \mathrm{E}_t \pi_{t+1} + \kappa x_t,$$

and the specification of the error term ε_t. Under either optimal discretion or commitment, $x_t = \pi_t = 0$ in equilibrium, and

$$i_t = \varepsilon_t = \left(\epsilon_{g,t} - \mathrm{E}_t \epsilon_{g,t+1}\right) - \left(z_t - \mathrm{E}_t z_{t+1}\right). \tag{12}$$

This equilibrium cannot be achieved with the policy rule (8). Thus, it may be that a LAW policy (i.e., $\theta_L > 0$) appears optimal only because the exercise forces one to find the best rule among a class of suboptimal rules. In this example, a cost-benefit trade-off arise only because policy is (suboptimally) restricted to follow a simple instrument rule.

Because an interest rate smoothing objective appears in the loss function (9), a zero output gap and zero inflation will not be the equilibrium, as it was in my simple example. But there is a general point to make. I would like to see optimal policy as the reference case so that one can assess the extent to which responding to credit in a simple rule is able to approximate a possibly more complex but optimal policy response to the possibility of credit crises. As it stands, any effects of the credit shock and a crisis on output and inflation arise only because (i) Δi_t^2 is in loss function and (ii) policy is suboptimal in following a simple rule.

5.3 Limitations of the Approach

The authors adopt a simple model that helps to organize and guide their analysis. But, as a simple model, it naturally has limitations. Let me mention three.

First, as in Ajello et al. (2016), there is no real role for credit in the model. Absent a crisis, the model displays a type of "credit dichotomy": the evolution of real activity and inflation is independent of credit growth (unless of course policy responds to credit growth). In models displaying the classic "monetary dichotomy," the real side was independent of money demand and supply, and only nominal values were affected by shifts in either the demand or supply of money. Here, the separation is more extreme. The specification of the equation for credit growth affects nothing except credit growth and the probability of a crisis. Rapid, possibly distortionary, credit growth does not affect investment, consumption, aggregate demand, marginal costs, inflation, or any other variable. This may be useful as a simplifying assumption, allowing the analysis to focus on the crisis implications of credit, but it does represent a limitation. For example, if financial markets are not frictionless in non-crisis times, excessive credit growth may reflect a misallocation of resources, and there may be a benefit of LAW policies that this analysis (and others) are missing. And the model cannot provide any guidance on whether excess credit growth creates distortions that would justify it appearing in the policy objective function along with inflation and a measure of real activity.

Second, the authors, following Ajello et al. (2016), assume that private agents do not see a crisis coming. As a description of the period leading up to the global financial crisis, that is a plausible assumption, though the crisis also caught policymakers by surprise. The assumption of myopic private agents is more problematic as the benchmark in an analysis of post-crisis monetary policy and may not be the best benchmark for analyzing LAW policies. It is not just central bankers who are more aware of the potential for financial crises; so are private agents. When private agents react to expectations of a future crisis, (12) shows that the optimal policy response is to raise the nominal interest rate, that is, to lean with the wind rather than against it.

Third, any effective lower bound on the policy interest rate is ignored. Taking this additional constraint into account might increase the value of LAW policies that reduce the probability of and the severity of a crisis.

6. Conclusions

Gerdrup et al. make a real contribution in providing an evaluation of LAW policies in the context of a dynamic, quantitative, open-economy model, using historical data on credit and crises to endogenize the probability of crises and their severity. Within the class of policy instrument rules they investigate, they find a clear role for LAW policies when private agents, but not the central bank, ignore the possibility of a crisis. The benefits in reducing the chance that a crisis occurs and the severity of one if it does occur outweigh the costs arising from a policy that keeps interest rates higher than would be justified by a focus solely on inflation and output.

Interesting extensions of their analysis would be to adopt a model that does not display a credit dichotomy but instead incorporates real linkages between credit growth and economic activity, to allow pre-crisis policies that dampen the severity of a crisis to also affect the probability of exiting a crisis, and to consider whether there are long-term effects such as scarring or growth effects of a financial crisis. Such costs may significantly affect the cost-benefit calculation of LAW policies.

References

Agur, I., and M. Demertzis. 2013. "Leaning against the Wind and the Timing of Monetary Policy." *Journal of International Money and Finance* 35 (June): 179–94.

Ajello, A., T. Laubach, D. Lopez-Salido, and T. Nakata. 2016. "Financial Stability and Optimal Interest-Rate Policy." Finance and Economics Discussion Series No. 2016-067, Board of Governors of the Federal Reserve System.

Anundsen, A. K., K. Gerdrup, F. Hansen, and K. Kragh-Sorensen. 2016. "Bubbles and Crises: The Role of House Prices and Credit." *Journal of Applied Econometrics* 31 (7): 1291–1311.

Filardo, A., and P. Rungcharoenkitkul. 2016. "A Quantitative Case for Leaning against the Wind." BIS Working Paper No. 594 (December).

Gambacorta, L., and F. M. Signoretti. 2014. "Should Monetary Policy Lean against the Wind? An Analysis Based on a DSGE Model with Banking." *Journal of Economic Dynamics and Control* 43 (418): 146–74.

Gelain, P., K. J. Lansing, and G. J. Natvik. 2015. "Leaning against the Credit Cycle." Working Paper No. 4/2015, Norges Bank.

Jordà, Ò., M. Schularick, and A. M. Taylor. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking* 45 (s2): 3–28.

———. 2017. "Macrofinancial History and the New Business Cycle Facts." In *NBER Macroeconomics Annual 2016*, Vol. 31, ed. M. Eichenbaum and J. A. Parker, 213–63 (chapter 3). University of Chicago Press.

Justiniano, A., and B. Preston. 2010. "Can Structural Small Open-Economy Models Account for the Influence of Foreign Disturbances?" *Journal of International Economics* 81 (1): 61–74.

Olsen, Ø. 2015. "Integrating Financial Stability and Monetary Policy Analysis." Speech at the Systemic Risk Centre, London School of Economics, London, April 27.

Svensson, L. E. O. 2016. "Cost-Benefit Analysis of Leaning against the Wind: Are Costs Larger Also with Less Effective Macroprudential Policy?" NBER Working Paper No. 21902.

———. 2017. "How Robust Is the Result That the Cost of 'Leaning Against the Wind' Exceeds the Benefit? Response to Adrian and Liang." (January).

Walsh, C. E. 2009. "Using Monetary Policy to Stabilize Economic Activity." In *Financial Stability and Macroeconomic Policy*, 245–96. Proceedings of a symposium sponsored by the Federal Reserve Bank of Kansas City, held in Jackson Hole, Wyoming, August 20–22.

Monetary Policy, Private Debt, and Financial Stability Risks*

Gregory H. Bauer[a] and Eleonora Granziera[b]
[a]Bank of Canada
[b]Bank of Finland

Can monetary policy be used to promote financial stability? We answer this question by estimating the impact of a monetary policy shock on private-sector leverage and the likelihood of a financial crisis. Impulse responses obtained from a panel VAR model of eighteen advanced countries suggest that the debt-to-GDP ratio rises in the short run following an unexpected tightening in monetary policy. As a consequence, the likelihood of a financial crisis increases, as estimated from a panel logit regression. However, in the long run, output recovers and higher borrowing costs discourage new lending, leading to a deleveraging of the private sector. A lower debt-to-GDP ratio in turn reduces the likelihood of a financial crisis. These results suggest that monetary policy can achieve a less risky financial system in the long run but could fuel financial instability in the short run. We also find that the ultimate effects of a monetary policy tightening on the probability of a financial crisis depend on the leverage of the private sector: the higher the initial value of the debt-to-GDP ratio, the more beneficial the monetary policy intervention in the long run, but the more destabilizing in the short run.

JEL Codes: E52, E58, C21, C23.

*We would like to thank Oscar Jordà, Jim Hamilton, Harrison Hong, Alexander Ueberfeldt, and seminar participants at the 2016 IJCB Annual Research Conference, Bank of Finland, Bank of Italy, Bank of Spain, Bundesbank, European Central Bank, and Universitat Pompeu Fabra for helpful comments. Graeme Westwood provided outstanding research assistance. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada or the Bank of Finland. Corresponding author: eleonora.granziera@gmail.com.

1. Introduction

In the wake of the recent financial crisis, policymakers have disagreed about whether central banks should extend their inflation-targeting mandates to promote financial stability. The supporters of "leaning against the wind" policies argue that low-for-long policy rates may help contribute to the buildup of debt that ultimately poses a risk to financial stability. This is particularly true in the housing market, where low rates may encourage households to take on larger mortgages and spur house price overvaluation (Taylor 2007, 2010). Because of these concerns, some have argued that monetary authorities should raise interest rates more than is warranted by contemporaneous price and output stability objectives alone (Borio 2014).

As monetary policy has a wide impact on both the economy and financial markets, setting interest rates higher than suggested by the central bank monetary policy rule might be costly in terms of lower inflation and lower output. For this reason, other policymakers believe that financial stability should be achieved separately through targeted financial regulation and supervision (e.g., Svensson 2014). Monetary policy should then be used at most to "clean"; i.e., policy should stabilize output and inflation after the bust in asset prices has occurred (Greenspan 2002, Mishkin 2011).

There is a middle ground: even if monetary policy is viewed as the last line of defense, there might be a case for monetary policy intervention to address financial stability concerns when large financial imbalances affect the short-term outlook for output and inflation. Moreover, when a low interest rate environment has encouraged the buildup of broadly based—as opposed to sector-specific—imbalances, monetary policy might complement macroprudential policy and directly contribute to financial stability.[1] However, this view relies on the efficacy of leaning when imbalances are higher than normal.

[1]See Smets (2014) for a detailed summary of the views on the role of monetary policy and macroprudential policies in promoting the soundness of the financial system.

In all of these cases, the rationale for central bank intervention rests on the assumptions that monetary policy is able to reverse the buildup of excess leverage and that "leaning against the wind" policies can be beneficial. However, the effect of monetary policy interventions on key risk indicators—such as the debt-to-GDP ratio—is a priori ambiguous: on the one hand, contractionary policies might lower both debt and the probability of a future crisis over the medium run, while on the other hand, they might have a negative impact on inflation and real activity in the short run. Thus, there are a number of outstanding questions facing policymakers. How effective is monetary policy in lowering the leverage of the private sector along with the subsequent likelihood of a financial crisis? Does this happen over the short or long run? Is leaning better undertaken when financial imbalances are low, or should it occur only when the likelihood of a crisis has become more apparent?

In this paper we contribute to the debate on the leaning versus cleaning roles of monetary policy by investigating whether monetary policy can achieve less leveraged and ultimately less risky economies. We structure our analysis in two steps. In the first step, we assess whether monetary policy can successfully decrease the leverage of the private sector, measured via the private debt-to-GDP ratio. As individual countries have experienced only a small number of financial cycles and financial crises during the last forty years, it would be difficult to determine the impact of a monetary policy shock on financial stability using data from a single country alone. For this reason we consider a broad cross-section of countries that have experienced financial cycles and crises of varying amplitude, size, and timing. In particular, we construct a comprehensive data set for eighteen advanced economies from 1975:Q1 to 2014:Q4. The data include measures of real economic activity, prices, credit, and financial variables. We obtain an aggregate impulse response function of the private debt-to-GDP ratio to a monetary policy shock by averaging impulse responses estimated from individual country VARs. In contrast to most of the existing studies, we use sign restrictions to identify a contractionary structural monetary policy shock while remaining agnostic about the response of the private debt-to-GDP ratio.

In the second step, we relate the private-sector leverage to the likelihood of a financial crisis occurring in the future by using a cross-country panel logit regression. Conditional on the response of the variables to the monetary policy shock and on the panel logit estimates, we can then evaluate the effects of the unanticipated tightening on the subsequent likelihood of a crisis. Some recent studies have analyzed the effect of monetary policy shocks on private debt, while others have focused on the determinants of financial crises. To the best of our knowledge, this is the first empirical paper that studies the effects of monetary policy on private debt and its repercussions for the likelihood of financial crises in a unified framework.

Our results provide new evidence on the questions posed above. A sizable, contractionary monetary policy shock causes real private debt (in deviation from trend) to rise on impact, as nominal debt barely responds and inflation falls. As output shrinks, the debt-to-GDP ratio increases in the short run. In the medium to long run, output recovers and debt decreases, ultimately resulting in a decline in the debt-to-GDP ratio. However, as the probability of a financial crisis increases with the debt-to-GDP ratio, we find that initially a crisis becomes more likely. A contractionary policy is thus risky in the short run. Eventually, as the cumulative response of the debt-to-GDP ratio turns negative, the probability of a crisis drops to a lower value than it was before the shock. Our analysis thus suggests the existence of an intertemporal trade-off that is new to the literature: while an unexpected monetary policy tightening might reduce leverage and promote financial stability in the long run, it might actually generate financial fragility in the short run. This inverse J-curve pattern indicates that additional macroprudential tools may be required or the economy may have to experience short-term pain to achieve long-term gain.

A further novel result is that the effectiveness of monetary policy in mitigating the probability of a financial crisis depends on the current leverage of the economy. In general, an unexpected tightening will ultimately be more effective the higher is the initial debt-to-GDP ratio with respect to trend, but in the short term, it will be more destabilizing. Waiting until financial imbalances are quite high before increasing policy rates can be risky. This result helps policymakers evaluate the trade-off between starting to lean sooner rather than later.

Our paper is related to two growing strands of the literature that discuss the role of monetary policy and financial stability. The first strand examines the role of monetary policy and private debt. From a theoretical perspective, Gelain, Lansing, and Natvik (2015) and Alpanda and Zubairy (2017) use a dynamic stochastic general equilibrium (DSGE) model to find that an unexpected tightening of policy rates affects real private debt through the housing sector by increasing mortgage rates and discouraging new lending. However, it also directly affects the non-housing sectors of the economy. In these models, only new loans respond to a monetary policy shock. As these loans represent a small fraction of total debt, the response of real debt is negative but moderate, implying an increase in the debt-to-GDP ratio in the short run. Svensson (2013) reaches the same conclusion using a partial equilibrium model of household debt dynamics calibrated to the Swedish economy. We include a national house price index in our model to help account for the importance of housing market dynamics. Our findings confirm empirically the results in these structural studies using a broad cross-section of country experience.

The existing empirical evidence on the effect of monetary policy on private debt is limited and focuses mostly on single-country estimates (e.g., Angeloni, Faia, and Lo Duca 2015 for the United States, Laséen and Strid 2013 for Sweden, and Robstad 2017 for Norway). Goodhart and Hofmann (2008) argue that estimates from individual countries might be imprecise, and they conduct a panel VAR analysis of house prices, money, and credit for a sample of industrialized economies. In all of these studies, real debt decreases after a contractionary monetary policy shock, but there is no consensus about the response of the debt-to-GDP ratio. Most of these studies use short-run zero restrictions to identify a monetary policy shock and constrain debt so as not to respond to a monetary policy shock on impact.

There is a growing interest in evaluating the role of private-sector leverage as a potential source of financial instability. Babecký et al. (2012), Gourinchas and Obstfeld (2012), Schularick and Taylor (2012), and Aikman et al. (2015) show that private domestic credit expansion is a significant predictor of financial crises. However, these papers do not investigate how monetary policy can affect credit. To the best of our knowledge, we provide the first empirical assessment of how monetary policy shocks translate into changes in leverage that affect the likelihood of a financial crisis.

Finally, a small number of recent papers use structural models to weigh the benefits and costs of “leaning against the wind” policies. Ajello et al. (2016) and Alpanda and Ueberfeldt (2016) introduce DSGE models for the United States and Canada, respectively, where the economy faces an endogenous probability of a financial crisis that is a function of private credit conditions. Svensson (2016) constructs an analytic model—calibrated to the Schularick and Taylor (2012) likelihood of a crisis—to show that leaning can be expensive, as it implies that when a crisis commences, the economy is already bearing a higher level of unemployment. Gerdrup et al. (this issue) construct a New Keynesian model in which the likelihood of the economy entering a financial crisis follows a Markov-switching process that depends on credit growth. Leaning is also found to be expensive, as it generates more volatile output and inflation in non-crisis periods. These calibrated studies have highlighted an intertemporal trade-off faced by monetary authorities: stabilization of output and inflation in normal times versus decreasing the probability and costs of a future financial crisis. We show new dynamics for the risk of contractionary policies and detail whether these policies should be implemented earlier or later in the cycle.

In interpreting our results, we stress that we are not conducting a welfare analysis to weigh the benefits and costs of monetary policy tightening. In addition, we do not provide a metric to weigh short-term losses against long-term gains. Furthermore, our identified monetary policy shock is not directly interpretable as a systematic “leaning against the wind” policy. These considerations can be addressed in the context of general equilibrium models such as in Ajello et al. (2015) and Alpanda and Ueberfeldt (2016). Therefore, we provide an answer to the question of whether monetary policy *could*, rather than *should*, be used to promote financial stability.

The paper is organized as follows: Section 2 provides some stylized facts regarding the evolution of private debt and the occurrence of financial crises. Section 3 describes the panel VAR model and shows the impulse responses to a monetary policy shock. Section 4 reports the effects of leverage dynamics on the probability of a crisis and evaluates the impact of an unanticipated monetary policy tightening on the likelihood of a crisis. Section 5 concludes.

2. Private Debt and Financial Crises: Stylized Facts from Cross-Country Evidence

In this paper, we consider private-sector leverage as an indicator of financial stability and show that it acts as an important channel through which monetary policy can affect the probability of a financial crisis. We start by detailing the data on debt and then turn to our definitions of financial crises.

2.1 Private Debt

We focus on leverage, as excessively leveraged economies might be less resilient to shocks and have lower loss-absorption capacities. Private-sector leverage is measured as the ratio of nominal private debt to nominal output. Private debt is defined as total credit provided by domestic banks to the domestic, private, non-financial sector.[2] We base our analysis on private credit provided by banks, as Jordà, Schularick, and Taylor (2016) stress that bank credit is the predominant form of private-sector borrowing in advanced economies. The data for all countries in our analysis are collected by the Bank for International Settlements (BIS) and cover the sample 1975:Q1 to 2014:Q4.[3]

Figure 1 shows the weighted average of the debt-to-income ratio across countries where weights are given by countries' shares of output.[4] In each quarter, the light grey area spans the range of values

[2]The private non-financial sector includes non-financial corporations, households, and non-profit institutions serving households. The series capture the outstanding amount of credit at the end of the reference quarter. Credit covers loans as well as debt securities.

[3]The countries included in our sample are Australia, Belgium, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, Netherlands, New Zealand, Norway, Spain, Sweden, Switzerland, the United Kingdom, and the United States. We consider the longest sample for which credit is available for all countries.

[4]We compare our data on credit with the aggregate private bank credit to the private non-financial sector as percentage of GDP shown in Jordà, Schularick, and Taylor (2016), available at the annual frequency from 1870. The series provided by Jordà, Schularick, and Taylor shows the same dynamics as the (unweighted) average of our series. The levels of both series are about 55 percent in 1975, and rise throughout the sample. They dip slightly to 80 percent in 1995, reach a peak of almost 120 percent in 2008–09, and plummet in 2010.

Figure 1. Private Debt-to-GDP Ratio, 1975:Q1–2014:Q4

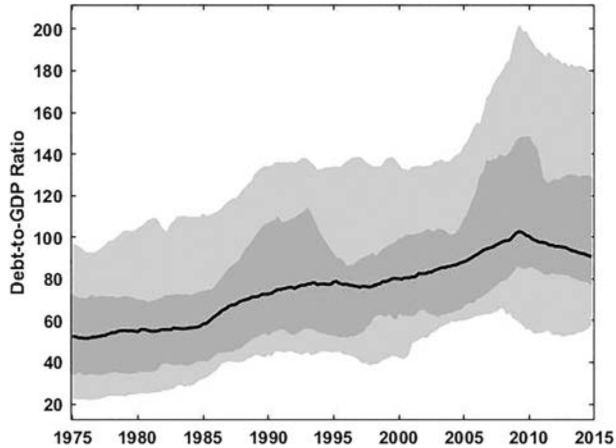

Notes: Private debt-to-GDP ratio for eighteen advanced economies over the sample 1975:Q1–2014:Q4. The black line shows the weighted average across countries, where countries are weighted by their share of output. The light grey shaded area delimits the range of values of the individual country ratio; the dark shaded area, the 20 percent and 80 percent percentiles.

(minimum and maximum) across countries, while the dark shaded area delimits the 20th and 80th percentiles.

Private debt as a fraction of GDP has risen rapidly and substantially across countries, increasing from an average of about 55 percent in 1975:Q1 to 90 percent in 2014:Q4. The strong rise in leverage can be attributed primarily to a substantial increase in mortgage credit, which represented one-third of bank assets at the beginning of the twentieth century and constitutes about two-thirds today (Jordà, Schularick, and Taylor 2015). At the same time, household borrowing was facilitated by higher levels of financial development (Chen et al. 2015). The level of the debt-to-GDP ratio shows a high degree of heterogeneity across countries which increased substantially during the financial crisis. The cross-sectional distribution of the debt-to-GDP ratio can be explained by different legal and economic institutions, including legal enforcement of contracts (e.g., the time needed to repossess a house (Bover et al. 2016)).

The upward trend in the debt-to-GDP ratio poses some challenges for empirical analysis. Moreover, the literature that studies the determinants of financial crises finds that the short-term dynamics of credit or the deviations of debt-to-GDP from trend, rather than the level, are significant predictors of crises (e.g., Babecký et al. 2012, Gourinchas and Obstfeld 2012, Schularick and Taylor 2012, and Drehmann 2013). We therefore use the debt-to-GDP gap, measured as deviation from a one-sided Hodrick-Prescott (HP) filter trend, in the analysis that follows. To account for the fact that credit cycles are characterized by longer duration and larger amplitude than those of traditional business cycles (Aikman et al. 2015, 2016), we use a much larger smoothing parameter ($\lambda = 400{,}000$) than the one commonly used for quarterly data. The resulting trend is thus very slow moving. This definition of the debt-to-GDP gap is also adopted under Basel III for the implementation of countercyclical capital buffers (Drehmann 2013).

While the trend is often ascribed to financial deepening, as financial innovations granted access to credit markets to previously unserved households and businesses, gaps in debt-to-GDP ratio may reflect several causes. For example, credit expansions may be driven by active risk-taking of financial intermediaries due to incentives that may not be fully aligned with those of shareholders (e.g., Allen and Gale 2000 and Bebchuk, Cohen, and Spamann 2010). Alternatively, shareholder risk appetite may be elevated (Danielsson, Shin, and Zigrand 2012; Adrian, Moench, and Shin 2013). Widespread optimism may be shared by financial intermediaries and other agents in the economy (Reinhart and Rogoff 2009; Gennaioli, Shleifer, and Vishny 2012, 2013; Barberis 2012). In these environments, banks might have an incentive to underwrite poor-quality loans and seek risk, posing threats to financial stability. In the analysis that follows, we implicitly interpret large positive deviations from trend as reflecting expansions in "bad" credit due to the incentives just described.

While there seems to be an upward trend in the ratio for all of the economies considered, the short-term dynamics of individual countries are quite heterogeneous.[5] Figure 2 shows the number of

[5]For example, while some countries—such as Ireland, Spain, and the United States—have gone through a large leveraging-deleveraging episode during the recent financial crisis, others—such as Canada, Switzerland, and Sweden—are still experiencing an increase in borrowing relative to income.

Figure 2. Financial Cycles, 1975:Q1–2014:Q4

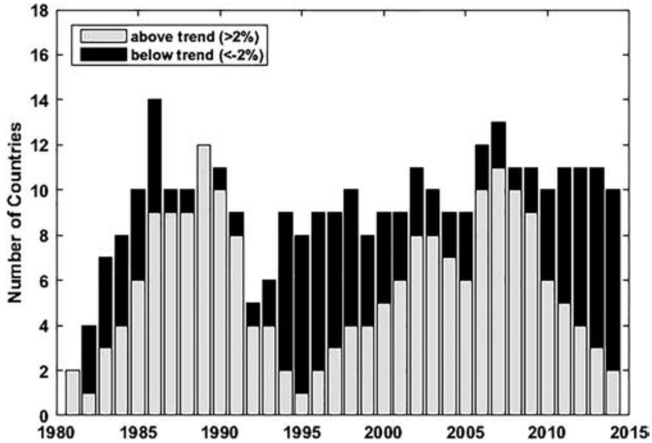

Notes: Number of countries with debt-to-GDP ratio on average at least 2 percent above (below) trend in grey (black) in a given year. The trend is constructed using a one-sided HP filter with smoothing parameter of 400,000.

countries in which the debt-to-GDP gap is higher or lower than 2 percent in a given year. The figure suggests that while cyclical fluctuations might be correlated, they are not perfectly synchronized.[6] For example, even during the recent financial crisis, when most countries experienced an increase in debt-to-GDP ratio well above trend, some countries were deleveraging.

Moreover, looking at individual cross-country correlations, in some cases debt-to-GDP gaps are strongly positively correlated—e.g., Finland and Sweden (0.83), Australia and New Zealand (0.86), or Ireland and Spain (0.88)—while in others, the gaps are negatively correlated, as in Japan and the Netherlands (–0.70) or Germany and Sweden (–0.79).

In our data, private-sector credit is strongly procyclical, as the average contemporaneous correlation between the growth rates of real output and real private debt is 50 percent. This confirms a similar result in Jordà, Schularick, and Taylor (2016).

[6]This finding is consistent with other studies; e.g., Baron and Xiong (2016).

The evolution of the housing market might be an important factor that can explain leveraging and deleveraging episodes. In fact, overvaluations in house prices might reflect high investor risk appetite and over-optimism, possible sources of large and rapid credit expansions. Indeed, in our sample, the contemporaneous correlation between debt-to-GDP ratio and real house prices ranges from 18 percent in Germany to 65 percent in the United States. Moreover, mortgage loans represent the majority of bank assets. Therefore, in the analysis that follows, we model the evolution of house prices and credit dynamics jointly.

2.2 Financial Crises

In our analysis, we construct two alternative measures of financial crises. In the first, we consider systemic banking crises that involve a large number of banks and therefore pose a threat to the entire economy. To identify the crises, we use the classification provided by Caprio and Klingebiel (2003) and Laeven and Valencia (2012). In data extending to 2003, Caprio and Klingebiel (2003) identify financial crises as periods of significant and systemwide financial distress in the banking system. In a longer data set, Laeven and Valencia (2012) impose an additional requirement that the distress must be followed by widespread insolvencies or significant banking policy interventions.[7] In our application, we combine the two data sets. However, even after the two classifications have been combined, this definition of financial crises delivers only a small number of episodes in our sample. Moreover, the crises are difficult to date accurately, and most episodes coincide with the latest financial crisis of 2008–11.

For these reasons, we consider a second measure of financial stability risk: large bank equity corrections. Following Baron and Xiong (2016), we define a large bank equity correction as a decrease in the realized excess return on the national bank equity index of at least 25 percent over one quarter, or of at least 35 percent over two quarters. This definition identifies episodes of distress in the banking sector that might not lead to a fully fledged systemic banking

[7] Significant banking policy interventions include, for example, extensive liquidity support, bank nationalization, asset purchase, deposit freezes, or bank holidays.

Table 1. Comparisons: Systemic Banking Crises and Large Equity Corrections

	Systemic Banking		Equity Corrections
	LV	**LV + CK**	**BX**
Frequency and Duration			
Proportion of Quarters in a Crisis	0.11	0.18	0.14
Number of Crisis Episodes	18	26	88
Average Number of Quarters in a Crisis	16	19	4
Losses			
Cumulative Output Change (%)	−2.9	−1.2	0.5
Excess Return on Bank Equity (%)	−92.8	−66.3	−36.1
Non-Crisis Average			
Cumulative Output Change (%)	2.6	2.7	2.6
Excess Return on Bank Equity (%)	7.9	8.9	11.6

Notes: LV refers to the Laeven and Valencia (2012) classification, LV + CK refers to an integrated data set from Laeven and Valencia (2012) and Caprio and Klingebiel (2003), and BX refers to the episodes identified as in Baron and Xiong (2016). Losses are computed on the first four quarters following the start of the crisis.

crisis. However, they clearly show an impaired financial sector and likely translate into a reduction of the supply of financial services to the private sector. A monetary policy shock that reduces debt and leverage might decrease the likelihood of these tail episodes.

Table 1 shows the frequency, duration, and losses associated with the two definitions of crises. We count twenty-six episodes of systemic banking crises occurring over the sample 1975:Q1–2011:Q4, of which ten occur during 2007:Q3–2011:Q4. The stricter Laeven and Valencia (2012) classification identifies eighteen crisis events; using this definition, most countries have experienced only one crisis, two have experienced two crises (Sweden and the United States), and three have never experienced any (Australia, Canada, and New Zealand). Systemic banking crises are accompanied by large output losses with a cumulative decline in GDP in the first year of the crisis of 2.9 percent on average for the Laeven and Valencia (2012) data and a decline of 1.2 percent for the integrated data set. This compares with GDP growth rates of 2.6 (2.7) percent per annum on average in non-crisis periods. Using the two combined data sets, the average duration of the crises is nineteen quarters.

Figure 3. Frequency of Systemic Banking Crises and Bank Equity Corrections, 1975–2011

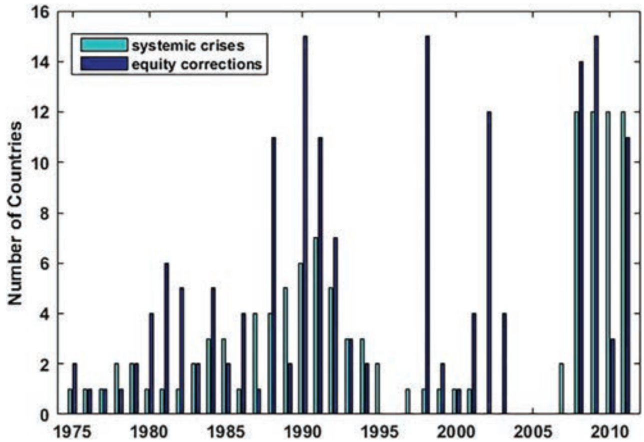

Notes: Number of countries in a systemic banking crisis or large bank equity correction in a given year. A country is counted as having experienced a crisis episode in a given year if it has experienced a crisis episode in at least one quarter of that year. The classification of systemic banking crisis is based on Caprio and Klingebiel (2003) and Laeven and Valencia (2012). We follow Baron and Xiong (2016) for large bank equity corrections.

Large bank equity corrections are more frequent than systemic banking crises, with eighty-eight occurrences of the former. However, they are shorter lived, lasting an average of only four quarters. They are also less severe disruptions to the financial system: while the excess rate of return on bank equity is –36.1 percent during the first year of a correction, this is smaller than the 92.8 percent and 66.3 percent declines recorded during the two types of systematic crises. GDP grows by a scant 0.5 percent during the first year of a bank equity correction compared with 2.6 percent growth during all non-crisis periods. Thus, bank equity corrections are far more numerous but less costly.

Figure 3 shows the frequency of systemic banking crises and large bank equity corrections for each year in our sample. Large bank equity corrections are more evenly distributed than systemic banking crises. In general, systemic banking crises are accompanied by large equity crashes, at least in the first year. We also observe episodes of large bank equity drops that were not followed by systemic banking crises, such as the crashes in the early 2000s, contemporaneous to the stock market downturn of 2002.

To provide additional insight, we conduct an event-study analysis in which we show the dynamics of bank equity returns, debt-to-GDP gaps, output gaps, and house price gaps during an event window that runs from eight quarters before to eight quarters after a crisis episode. Figure 4 shows the average of variables around a systemic banking crisis (panel A), and a bank equity correction (panel B). The qualitative patterns of the variables are very similar across the two types of crisis episodes. Bank equity returns fall with the beginning of the crisis, but the largest drop occurs after five quarters in the case of systemic banking crises. Returns are both lower and more volatile following the crisis. In the run up to a crisis, debt-to-GDP gaps are positive and large, especially for systemic banking crises. While we observe some deleveraging right before the onset of the crisis, the speed of the adjustment in the debt-to-GDP ratio increases in the aftermath of the crisis. Both types of crises are associated with a significant decline in the output gap, which fails to close within the first two years. However, the recovery is more sluggish for systemic banking crises. As in the case of debt-to-GDP gaps, house prices are higher than trend before the crisis, but they start to decline two to three quarters before the onset of both types of crises and continue to correct after that. While the dynamics are qualitatively similar across crises, the magnitude of the drops in output gaps, bank equity returns, and house prices suggests that systemic banking crises are more costly than large bank equity corrections.

3. Monetary Policy and Financial Cycles

3.1 VAR Model

We conduct a cross-country analysis that yields an estimated average response of the debt-to-GDP ratio to a typical monetary policy shock across countries and across time. Rather than estimating a panel VAR which assumes that countries are "similar enough" to pool their credit cycles, we estimate individual country VARs and pool the individual impulse responses. We thus do not impose

Figure 4. Event Study

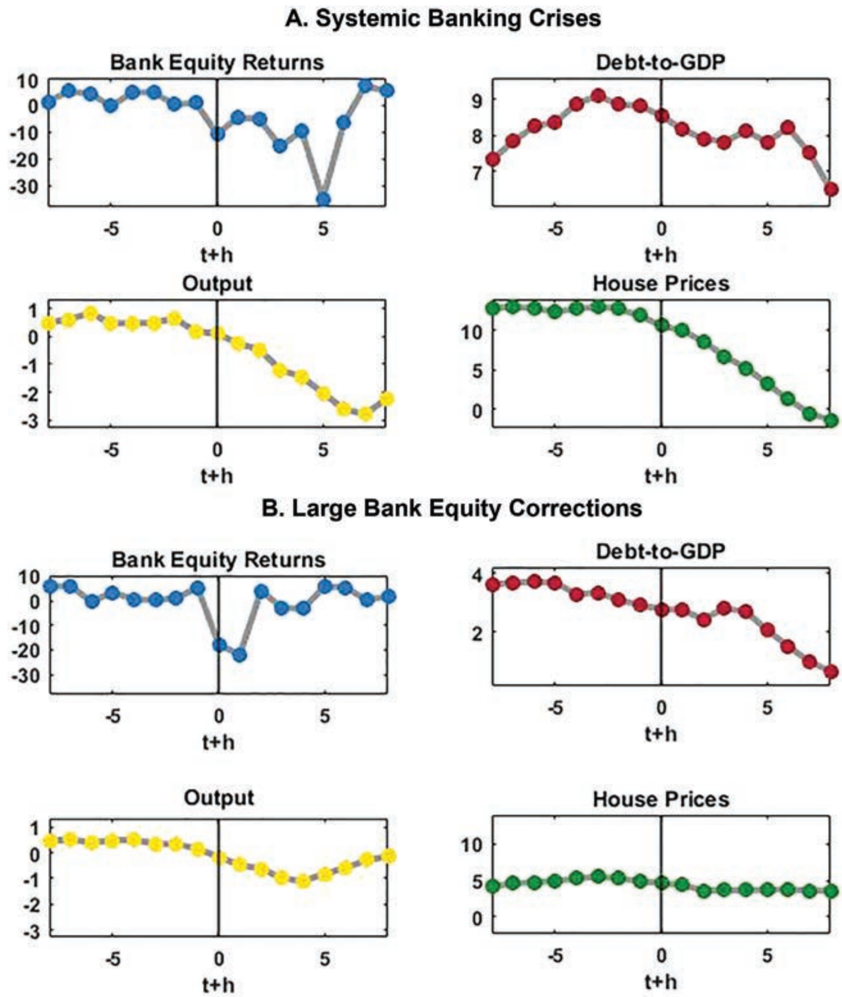

Note: Average of selected variables around crisis episodes: systemic banking crises (panel A); large bank equity corrections (panel B).

dynamic homogeneity, which can deliver biased estimates of the autoregressive coefficients in dynamic panel models even when T is large (Pesaran and Smith 1995). Dynamic heterogeneity might be induced by cross-country variation in the degree of financialization and other idiosyncratic characteristics that affect the transmission mechanism of monetary policy.

The eighteen country-i VARs each take the form

$$Y_{i,t} = A_i + \sum_{l=1}^{p} B_{i,l} Y_{i,t-l} + \varepsilon_{i,t} \qquad i = 1, \ldots, 18,$$

where $Y_{i,t}$ is a $K \times 1$ vector of endogenous variables, A_i is a $K \times 1$ vector of country-specific intercepts, $B_{i,l}$ is a $K \times K$ matrix of autoregressive coefficients, and $\varepsilon_{i,t}$ is a $K \times 1$ vector of errors with $\varepsilon_{i,t} \sim N(0, \Sigma_i)$. The vector of endogenous variables includes real output, inflation, a short-term interest rate, the real national house price index, and the private debt-to-GDP ratio. Real variables are obtained by deflating nominal variables by the CPI (all items). Output, house prices, and the debt-to-GDP ratio are expressed in percentage deviation from the trend calculated using a one-sided HP filter. Inflation is the quarter-over-quarter log-difference in CPI (all items). The short-term interest rate is the three-month Treasury-bill (T-bill) rate.

We estimate the model by least squares with a lag length of two over the sample period 1975:Q1–2014:Q4. A monetary policy shock is identified via sign restrictions as in Uhlig (2005). We impose the restriction that after a contractionary monetary policy shock (an increase in the short-term interest rate, normalized to a 100 basis point response on impact), both the output gap and inflation do not increase. In order to capture the effects in the housing market, the response of property prices is also restricted to be non-positive. The restriction imposed on real house prices is consistent with DSGE models that incorporate a housing sector such as Iacoviello (2005) or Iacoviello and Neri (2010) and the semi-structural model of Svensson (2013). Responses of all variables are constrained only on impact.

We remain agnostic about the behavior of the debt-to-GDP ratio by not restricting its response. For the numerator, an increase in the interest rate will likely discourage lending by increasing borrowing costs. However, a higher policy rate will also lower the inflation rate, further increasing the real cost of borrowing on the one hand, but increasing the real value of the existing debt stock on the other. Therefore, the total effect of a monetary policy shock on real debt will depend on how monetary policy interventions transmit to both mortgage rates and inflation. For the denominator, the shock will cause output to decline. Thus, the response of the debt-to-GDP ratio is a priori unclear and depends on the relative size of the responses of debt and output.

In the case of sign restrictions, inference is complicated by the fact that the impulse responses are not point identified, but only set identified; i.e., impulse responses are bounded up to an interval (Moon and Schorfheide 2012). While most of the existing literature uses Bayesian methods to derive the highest posterior density intervals in these models, we conduct inference using the novel frequentist procedure outlined in Moon, Schorfheide, and Granziera (2013), which is based on a Bonferroni approach. We prefer the frequentist approach, as the Bayesian approach for set-identified models is sensitive to the choice of priors even in large samples (Moon and Schorfheide 2012, Giacomini and Kitagawa 2015). Average impulse responses across countries are obtained by pooling individual country responses with weights corresponding to the country share of output.

3.2 Impulse Responses

Figure 5 shows the impulse response functions of each variable after an unanticipated 100 basis point increase in the short-term rate. We show the 68 percent confidence interval for the impulse responses as well as the average of the confidence set.

After a sizable, contractionary monetary policy shock, output falls below potential for about ten quarters. Inflation also falls, but by less than output, and the response becomes insignificant after one year. The private debt-to-GDP ratio increases on impact, although the response is only slightly significantly different from zero. The ratio falls below trend in the medium and long run.

To help with the interpretation of this latter result, we also compute the impulse responses of the real private debt and nominal private debt in deviation from a one-sided HP-filter trend. The responses are obtained from analogous structural VAR models where the real or nominal debt is used in place of the private debt-to-GDP ratio. These additional exercises, shown in figure 6, reveal that the monetary policy shock causes only a moderate reduction in nominal debt on impact. Similar to the results in Svensson (2013)

Figure 5. Responses to a Monetary Policy Shock

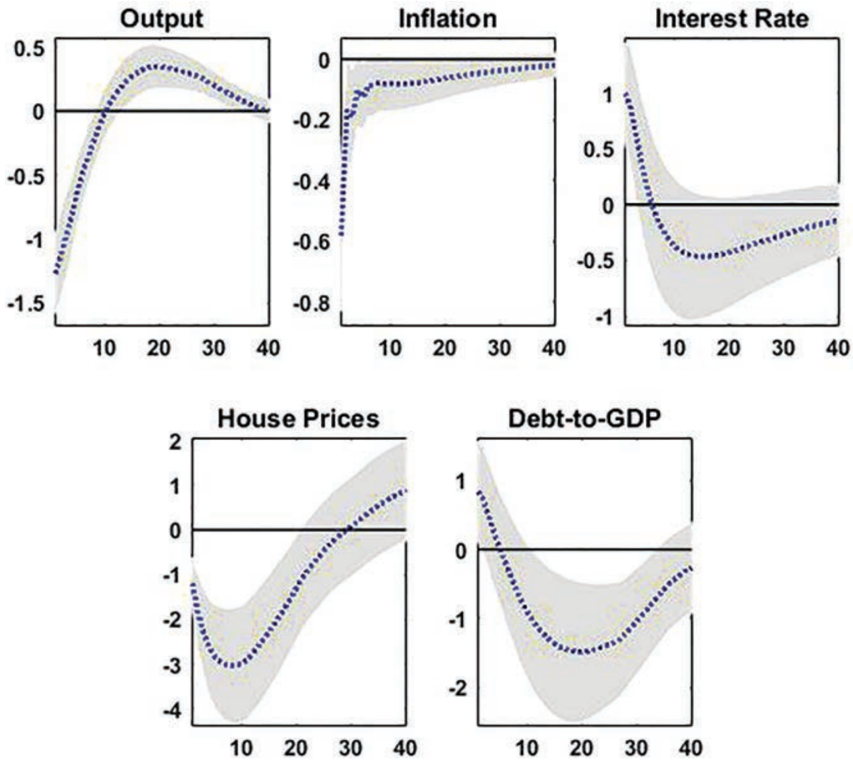

Notes: Impulse response functions after a 100 basis point contractionary monetary policy shock. Dashed lines indicate the average response; shaded areas indicate 68 percent confidence set obtained with the frequentist procedure in Moon, Schorfheide, and Granziera (2013).

and Gelain, Lansing, and Natvik (2015), the stock of nominal debt exhibits considerable inertia, as agents find it difficult to change existing contracts. Given that the fall in inflation is larger than the reduction in nominal debt, after a monetary policy shock real debt rises on impact. As nominal debt further decreases and inflation quickly rebounds, real debt falls below trend from the first quarter after the shock.

As output shrinks on impact, the debt-to-GDP ratio rises by about 0.85 percent (figure 5). However, starting from the first

Figure 6. Responses to a Monetary Policy Shock: Nominal Debt and Real Debt

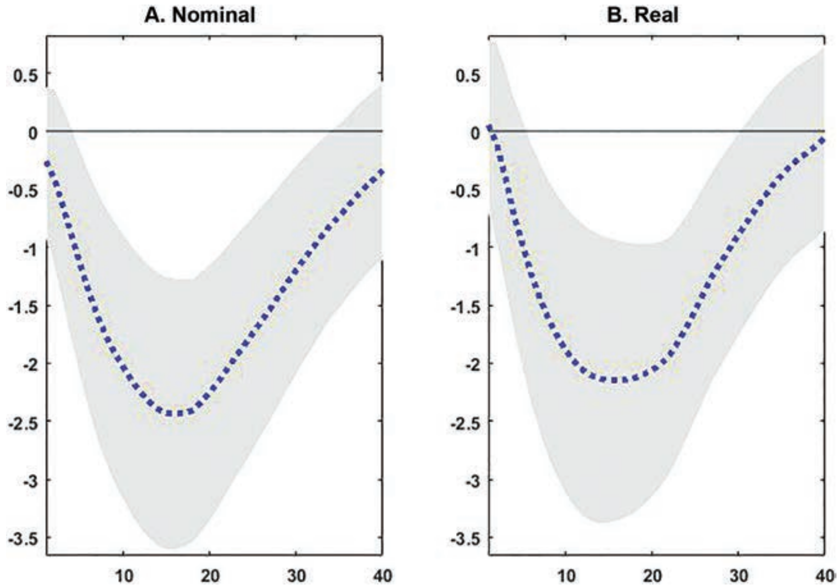

Notes: Impulse response functions of nominal debt (panel A) and real debt (panel B) after a 100 basis point contractionary monetary policy shock. Dashed lines indicate the average response; shaded areas indicate 68 percent confidence sets.

quarter after the shock, debt decreases and the output gap becomes smaller (in absolute terms). As the decline in the debt is larger than the fall in output, a tightening of monetary policy will decrease the debt-to-GDP ratio in the medium run and long run, starting after approximately six quarters. After five years from the initial tightening, the debt-to-GDP ratio is about 14 percent lower in percentage deviation from trend than before the tightening.

The effect of the monetary policy shock is more pronounced on house prices than it is on either real output or real debt. While house prices and output respond by roughly the same magnitude on impact, after two quarters the response of house prices is more than twice as strong as that of output. This result is consistent with other structural panel VAR studies such as Assenmacher-Wesche and

Gerlach (2008). House prices dip by about 3 percent one year after the shock and recover only in the long run.

Our results can be compared with a number of others from the literature. Using Norwegian data that runs from the mid-1990s to 2013, Robstad (2017) finds that the debt-to-GDP ratio rises following a contractionary monetary policy shock. However, Laséen and Strid (2013) conduct a similar analysis using Swedish data and report a decline in the debt-to-GDP ratio after a tightening. Using a pre-financial crisis sample of U.S. data, Angeloni, Faia, and Lo Duca (2015) find that both industrial production and household debt decline over all horizons considered. Although they do not provide the response of the debt-to-GDP ratio, the relative changes in the household debt and industrial production suggest that the debt-to-GDP ratio increases slightly in the short run but falls, although only moderately, in the long run. All of these studies are based on single-country evidence and restrict real debt not to change on impact in response to a contractionary monetary policy shock.

Our results are more comparable with those of Goodhart and Hofmann (2008), who estimate a panel VAR on seventeen countries over the 1970–2006 period. They do not restrict the contemporaneous response of debt. As in our analysis, after a 100 basis point increase in the interest rate, real debt and the debt-to-GDP increase on impact. However, while both debt and output decline from the second quarter after the shock, the response of debt is much stronger than the response of output. This leads to a substantial decline in the debt-to-GDP ratio in the medium and long run. The response of the debt-to-GDP ratio in our paper is qualitatively consistent with the structural studies of Svensson (2013), Gelain, Lansing, and Natvik (2015), Chen and Columba (2016), and Alpanda and Zubairy (2017).

3.3 Robustness Analysis

In this subsection we present a series of robustness checks on the VAR model. The characteristics of residential mortgage markets may affect the transmission of monetary policy (Assenmacher-Wesche and Gerlach 2008; Garriga, Kydland, and Sustek 2013). In particular, the proportion of fixed versus variable mortgage rates matters for the response of house prices, residential investment, and output to a monetary policy shock. We therefore investigate whether the responses of debt gaps and debt-to-GDP gaps differ according to the mortgage scheme by estimating the panel VAR on two subsamples of countries which are grouped according to their prevailing mortgage structure (Tsatsaronis and Zhu 2004). Panel A of figure 7 shows that the effect of a monetary policy tightening on the debt-to-GDP ratio is qualitatively and quantitatively very similar across the two samples.

We have examined a number of alternative specifications of the model. In the benchmark specification we prefer using the three-month rate rather than the target rate to mitigate the issues associated with the zero-lower-bound constraint, binding at the end of our sample. In a second experiment, we repeated our estimation with the policy rate in place of the three-month rate. Long-term interest rates are relevant for debt dynamics, as they will affect mortgage subscriptions. Hence, in a third exercise we add long rates to the variables in our baseline VAR specification. As the transmission of monetary policy might have been distorted during the financial crisis, in a fourth experiment we limit our estimation sample to 2006:Q4.

The results of these additional exercises are shown in panel B of figure 7. The impulse response functions are generally well within the confidence bounds of the baseline model throughout all horizons considered.

Despite its long-standing tradition in macroeconomics, the HP filter presents some drawbacks. Specifically, it generates spurious dynamic relations and is subject to the end-of-sample problem, i.e., the filter produces a trend component that is close to the observed data at the beginning and at the end of the sample. For these reasons, we have tried using data detrended by taking log-differences for output, house prices, and the debt-to-GDP ratio. We find that on impact the debt-to-GDP ratio in log-differences rises by more than the debt-to-GDP gap. However, the response for log-differences quickly converges to zero, as it is less persistent than the response of the gap. Log-differenced data would imply a permanent effect of monetary policy shocks on the level of the variables. Therefore we prefer working with HP-filtered series.

As further checks on the robustness of our results, we tried different identification strategies of the monetary policy shock. First,

Figure 7. Impulse Response Functions: Robustness Analysis

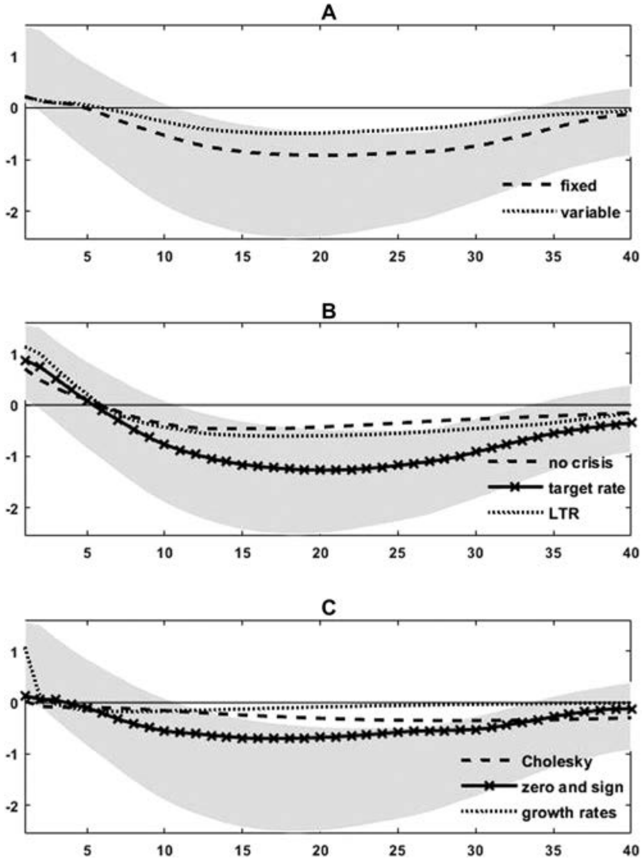

Notes: Impulse response functions of debt-to-GDP in deviation from trend after a 100 basis point contractionary monetary policy shock. Dashed lines indicate the average responses from different robustness experiments; shaded areas indicate the 68 percent confidence set for the impulse responses obtained from the baseline VAR.

we added a zero restriction on impact for output to the sign restrictions imposed on the other variables. While the qualitative response did not change, imposing a zero response on output dampens the impact of the monetary policy shock on the debt-to-GDP gap, as shown in panel C of figure 7.

Finally, we try to identify the monetary policy shock, resorting to short-run restrictions. In a first experiment where house prices and debt-to-GDP are ordered after the interest rate, we find qualitatively similar results to the ones obtained using sign restrictions for house prices, and the debt-to-GDP ratio, although the response is muted (panel C of figure 7). In a further experiment where the interest rate is ordered last and therefore debt cannot respond to monetary policy shocks on impact, we find that, differently from the sign-restriction identification strategy, in the short run debt-to-GDP ratio does not increase, while the medium- and long-run response is analogous to the one generated via sign restrictions.

We conclude that our specification is robust to a number of alternative scenarios.

4. Leverage Dynamics and Financial Stability Risks

We have shown that a monetary policy tightening decreases the debt-to-GDP gap, at least in the medium to long run. Although the magnitude of the decrease seems moderate, to evaluate its importance we should understand how the dynamics of the ratio affect the soundness of the financial system. The recent literature finds that a rapid expansion in credit is associated with a higher likelihood of experiencing a systemic banking crisis or a financial crisis recession (Gourinchas and Obstfeld 2012; Jordà, Schularick, and Taylor 2015). Baron and Xiong (2016) show that rapid credit growth may also lead to a large correction in bank equity prices.

The mechanism at work is the following: a large and rapid expansion of bank credit might drive funds to private-sector borrowers with poor credit quality, exposing banks to a higher number of defaults in case of adverse shocks. Such shocks may lead to a sharp drop in banks' equity prices. Because of the leading role played by banks in providing credit to the private sector, a sizable correction in the bank equity index might cause a decline in the supply of credit as leverage constraints become binding. Large drops in equity prices and the consequent depletion of bank capital might then trigger a systemic banking crisis which might require government interventions.

In this section we evaluate how reductions in leverage will translate into a less risky financial system by quantitatively assessing the linkages between deviations of debt-to-GDP from trend and the likelihood of a banking crisis.

4.1 Assessing the Likelihood of a Crisis

In order to understand the role of excessive credit as a source of financial instability, we estimate its impact on the probability of a systemic banking crisis or a large bank equity correction occurring in the near future using a panel logit regression framework. We follow Schularick and Taylor (2012) in using a cross-country panel, as the number of episodes for each country is limited and inference conducted on a single country may be misleading.

Following Gourinchas and Obstfeld (2012) and others, we construct a forward-looking dummy variable for crisis events as follows. Define $f_{i,t}^{k}$ to be the dummy variable that takes the value of 1 if a financial stability risk episode of type k starts within the next eight quarters in country i. The variable $k = \{sb, ec\}$ indicates the two types of crises that were detailed in section 2.2, where "sb" denotes a systemic banking crisis and "ec" an equity correction. We consider a horizon of eight quarters, which is of interest to monetary policy authorities. To reflect the uncertainty in the dating of the crisis episodes, we let $f_{i,t}^{sb}$ take the value of 1 in the first year of the crisis also.

Thus, our panel logit regression takes the form

$$\Pr\left(f_{i,t}^{k} = 1 \mid \delta_i^k, x_{i,t}\right) = \frac{\exp\left(\delta_i^k + \beta_k x_{i,t}\right)}{1 + \exp\left(\delta_i^k + \beta_k x_{i,t}\right)},$$

where $x_{i,t}$ is the set of explanatory variables. All specifications include country-specific fixed effects (δ_i^k) to allow for cross-country, time-invariant heterogeneity. Note that we run two separate regressions: in the first one the endogenous variable is a dummy indicating an upcoming systemic banking crisis $(k = sb)$, while in the second it is a dummy indicating a future bank equity index correction $(k = ec)$. Further, for both of these crisis-type events, we consider two regression models: one that includes only the debt-to-GDP ratio in deviation from trend, and a second one that includes all of the variables used in the VAR. By including only the debt-to-GDP ratio in the first regression, we impose that the other

Table 2. Panel Logit Regressions

	Systemic Banking Crises		Bank Equity Corrections	
	(1)	**(2)**	**(3)**	**(4)**
$debt/gdp_{t-1}$	0.128***	0.128***	0.027**	0.031**
	(0.0359)	(0.0410)	(0.0136)	(0.0144)
gdp_{t-1}		0.0136		0.145*
		(0.127)		(0.076)
inf_{t-1}		0.0575		−0.324**
		(0.241)		(0.127)
i_{t-1}		0.0493		−0.034
		(0.0725)		(0.0344)
rhp_{t-1}		0.0287		0.003
		(0.0207)		(0.009)
Observations	2,324	2,284	2,455	2,415
R^2	0.167	0.223	0.014	0.064
χ^2	12.626	18.312	4.089	19.313
p-value	0.000	0.003	0.043	0.002
Hit Rate	0.650	0.658	0.589	0.593
Standard Error	0.007	0.007	0.009	0.016
Quadratic Probability Score	0.163	0.153	0.357	0.344
Standard Error	0.015	0.016	0.002	0.017
AUROC	0.715	0.757	0.624	0.678
Standard Error	0.018	0.017	0.012	0.012

Notes: Estimates of a panel logit with country fixed effects. The dependent variable is the occurrence of a systemic banking crisis (columns 1–2) or large equity corrections (columns 3–4) between horizon $t + 1$ and horizon $t + 8$ quarters. "Systemic banking crises" refers to the integrated LV + CK data set, while "bank equity corrections" is defined similarly to Baron and Xiong (2016). Debt-to-GDP ratio is expressed in percentage deviation from an HP-filtered trend. Significance levels at 10, 5, and 1 percent are denoted as *, **, and ***, respectively. R^2 refers to the pseudo R-squared. Robust standard errors are obtained using the multiway clustering (country and time) of Cameron, Gelbach, and Miller (2011).

variables affect the probability of a crisis only through their effects on leverage.

Table 2 shows the results of the panel logit model for the two types of crises and the two sets of regressors. We focus on systematic banking crises first. The debt-to-GDP ratio is a positive, highly significant predictor of systemic banking crises (table 2, column 1). As the other variables from the VAR are not statistically significant predictors (column 2), we will use the debt-to-GDP ratio as the sole predictor in the subsequent analysis (section 4.2).

Figure 8. Financial Crisis Predicted Probabilities

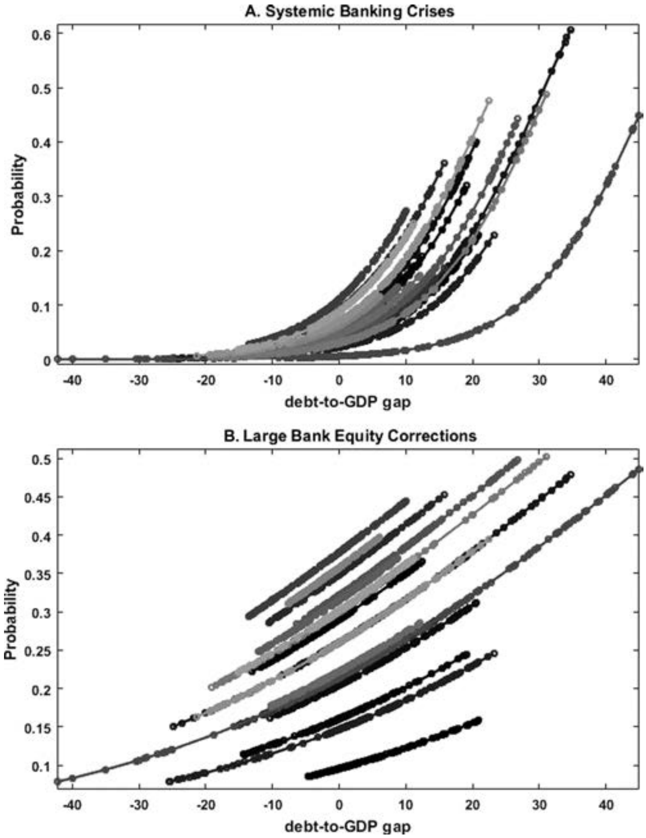

Notes: Probabilities are obtained from a panel logit regression in which the dependent variable is a dummy taking the value of 1 if a systemic banking crisis (panel A) or a large bank equity correction (panel B) occurs within the next eight quarters. The logit models, described in detail in section 4, include only the country fixed effect and the lagged debt-to-GDP ratio.

It is important to assess the fitted value of the model to understand the adjustments to the likelihoods that result from monetary policy innovations. Panel A of figure 8 shows the fitted values of the model for systemic banking crises across the countries in the sample, over the range of debt-to-GDP ratios experienced in sample. The graph shows that when the debt-to-GDP ratio is at its trend level, the estimated probability of a crisis starting sometime in the next two years ranges from approximately 1 to 12 percent. When the ratio is below this value, the probabilities compress to low levels. What is more interesting is the sharp increase in the likelihood of the onset of a financial crisis when the debt-to-GDP ratio moves above trend. For example, when the debt-to-GDP ratio is 20 percent above trend, the likelihood of a financial crisis for some countries can be quite high, at 30 percent or higher. At larger deviations of the ratio, the likelihoods can increase to 60 percent.

A similar exercise for the model to predict large bank equity corrections can be undertaken. The coefficient on the debt-to-GDP ratio is again positive and statistically significant (table 2, column 3). When the other variables from the VAR are added, we obtain significant coefficients on the inflation and output gap measures. However, the fitted values from the two specifications are similar, so we focus on the model containing only the debt-to-GDP ratio. It is interesting to note that the fitted values of this model display a more linear relationship between the debt-to-GDP ratio and the likelihood of an equity correction (figure 7, panel B). However, we note that the estimated levels may be quite high for this model as well. For example, when the ratio is at its trend value, the likelihood of a correction ranges from approximately 5 percent up to 35 percent. When the ratio is 20 percent above trend level, the likelihood may increase to 15 to 45 percent.

The ability of the debt-to-GDP ratio to forecast both systemic financial crises and bank equity corrections can be assessed using the area under the receiver operating characteristic curve (AUROC). This statistic estimates the likelihood of making correct decisions (Schularick and Taylor 2012, Jordà 2014). The area varies between 50 and 100 percent. In the results here the values of the AUROC statistic for models that contain only the debt-to-GDP ratio are 62.4 and 71.5 percent, suggesting that the ratio is more successful in predicting financial crises and corrections than is a random coin toss. We note that the values of these statistics are in line with those in Schularick and Taylor (2012), where the model to predict financial crises using the growth of private credit has an AUROC statistic of 67 percent.

4.2 Monetary Policy and Financial Stability Risks

In the previous sections, we completed two separate exercises. In the first, we derived the response of the leverage of the private sector (i.e., the debt-to-GDP ratio) to an unexpected monetary policy tightening. Then, we showed that the ratio, in deviation from trend, is a significant predictor of systemic banking crises and large bank equity corrections.

We are now in a position to answer the question as to how monetary policy can ultimately impact financial stability. We assess how an unanticipated monetary policy tightening affects the probability of a financial crisis through its effect on the debt-to-GDP ratio. The estimated impulse response of the debt-to-GDP ratio to a tightening can be calculated from the VAR model. The panel logit regression tells us how this affects the probability of a crisis. Given these two estimates, we compute the changes in the probability of a financial crisis that can be attributed to the unanticipated monetary policy shock.

The experiment is conducted as follows. We first compute the probability of a financial crisis occurring within the next two years for an initial, given deviation of the debt-to-GDP ratio from trend. This probability is obtained from the estimated panel logit model in table 2 that includes only the country fixed effects and the debt-to-GDP gap as predictors.[8] Next, given the cumulative responses of the debt-to-GDP ratio (in deviation from trend),[9] we can compute the value of the debt-to-GDP gap after the shock and the corresponding probability of a financial crisis. Finally, by comparing the probability of a financial crisis before and after the shock, we can assess whether the tightening has resulted in an increase or a decrease in financial stability risks.

For example, if we assume that the debt-to-GDP ratio is at trend before the shock hits, the results of the panel logit

[8]In this exercise, for all the initial debt-to-GDP gap values considered, we set the country fixed-effect parameters to their cross-country mean value. We are thus analyzing the effect of the tightening on the financial stability of a "typical" country.

[9]In the case of sign-restricted SVARs, the impulse response can be bounded up to an interval, rather than a point (Moon and Schorfheide 2012). We base our analysis on the mean of the identified interval.

regression imply that the probability of a systemic banking crisis (for an average country) is about 6 percent. The impulse response functions in figure 5 indicate that after a monetary policy shock the debt-to-GDP ratio increases by 0.85 percent, which results in an increase in the likelihood of a financial crisis to about 6.6 percent.

This exercise can be repeated for all of the horizons considered from the panel VAR. In addition, as the panel logit model is non-linear, the effect on the probability of a crisis will depend on the initial level of the debt-to-GDP ratio. To illustrate the non-linearities of the model, we repeat our experiment for different initial levels of the ratio. We thus assume that the debt-to-GDP ratio is at trend, or 5 and 10 percent above trend, before the monetary policy shock occurs. The results will indicate the differences among the effectiveness of tightening at different stages of financial fragility.

Figure 9 shows the probability of a systemic banking crisis (panel A) occurring within the next two years following the unexpected policy tightening. For all initial debt-to-GDP levels considered, the probability of a financial crisis rises after the shock, reaching a maximum increase approximately one year from the tightening. Then, as the debt-to-GDP ratio falls, the probability decreases sharply. As the cumulative effect on the ratio eventually becomes negative, the likelihood of a crisis falls below its initial value. Thus, five years after the shock, the likelihood of a financial crisis hitting the economy is considerably lower than it was before the shock. The magnitudes of the decreases in risk are quite large.

For example, if the initial debt-to-GDP ratio is 10 percent above trend, the likelihood of a financial crisis is approximately 18 percent (figure 9, panel A). Following the monetary policy shock, the decline in the debt-to-GDP ratio causes a decline of approximately 15 percent in the likelihood of a systemic financial crisis. For lower initial levels of the debt-to-GDP ratio, the declines are smaller. Indeed, regardless of the initial debt level, the likelihood of a crisis falls to between 0 and 5 percent, approximately, at the end of five years. Thus, the monetary policy tightening is able to reduce large amounts of tail risk.

A similar inverse J-curve pattern holds for the effect of a tightening on the likelihood of a large bank equity correction (figure 9, panel B). As these corrections are more common, the initial values of the likelihoods are higher than are their counterparts for

Figure 9. Effect of a Monetary Policy Shock on Crisis Probabilities

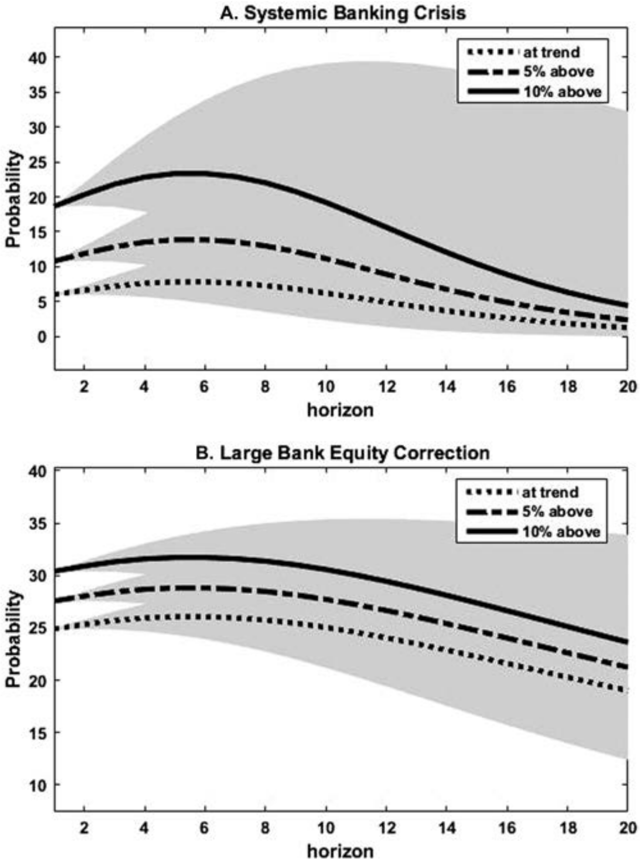

Notes: Probability of a crisis occurring within the next two years given the responses of the debt-to-GDP ratio, following a 100 basis point monetary policy shock. Probability of a financial crisis is calculated assuming that before the monetary policy shock hits, the debt-to-GDP is (i) at trend, (ii) 5 percent above trend, or (iii) 10 percent above trend. Shaded areas delimit the 68 percent confidence set.

full-fledged financial crises. For example, the likelihood of a large decline in bank equity values starting sometime in the next two years is approximately 31 percent when the debt-to-GDP ratio is 10 percent above trend. After the monetary policy shock, the likelihood declines rapidly, falling to about 24 percent after five years. We note that declines of similar magnitudes are experienced by economies where the initial debt-to-GDP levels are either at trend or 5 percent above trend (figure 9, panel B). However, the terminal values of the probabilities remain higher than for their systemic banking crisis counterparts. The tightening is thus able to substantially reduce the likelihood of an extreme left-tail event (i.e., a fully fledged financial crisis), while it has a much smaller effect on a more moderate episode (i.e., a large decline in bank equity).

4.3 Robustness Analysis

In this section we estimate some alternative specifications to show that our results are robust.

First, we estimate the effect of the debt-to-GDP gap on the probability of experiencing an episode of distress in the financial system, while taking into account that systemic banking crises can be interpreted as more severe events than banking equity corrections. In fact, table 1 and the upper panel of figure 4 show that systemic banking crises are accompanied by large declines in equity indexes. As the two types of crises have a natural ordering, we can estimate an ordered logit regression model that captures the likelihood of either type of crisis simultaneously. We thus construct a dummy variable that takes the value of 1 if a bank equity correction occurs and 2 if a systemic banking takes place. For quarters in which we experience both crises simultaneously, we assign the value 2 to the dummy. Table 3, columns 1–2, shows the results. As in the previous logit regressions, we find that the debt-to-GDP gap is a highly significant and positive predictor of crises. The variable retains its significance when a larger set of regressors is included in the logit model. The coefficient associated with the debt-to-GDP gap is between the value estimated in the logit regression for systemic banking crises and the one estimated for bank equity corrections.

Second, we also estimate the effect of the debt-to-GDP gap on the probability of experiencing either type of crisis; i.e., we investigate the ability of the debt-to-GDP ratio to simultaneously predict both financial crises and large bank equity corrections. In this second exercise we do not distinguish between the two types of financial stability risks based on their severity. We thus run a logit model where the dummy variable takes the value of 1 if we will experience either a

Table 3. Robustness: Logit Regressions

	Ordered Logit		**Logit-Pooled Crises Episodes**	
	(1)	**(2)**	**(3)**	**(4)**
$debt/gdp_{t-1}$	0.0508***	0.0538***	0.0487***	0.0483***
	(0.016)	(0.014)	(0.017)	(0.016)
gdp_{t-1}		0.117		0.142
		(0.085)		(0.107)
inf_{t-1}		−0.030		−0.038
		(0.106)		(0.127)
i_{t-1}		0.088**		0.149***
		(0.037)		(0.044)
rhp_{t-1}		0.002		0.016
		(0.007)		(0.010)
Observations	2,478	2,438	2,478	2,438
R^2	0.0255	0.0521	0.059	0.127
χ^2	217	357	566	848
p-value	0.000	0.000	0.000	0.000

Notes: Estimates of an ordered logit with country fixed effects (columns 1–2) and a logit model that pools the crisis episodes (columns 3–4). The dependent variable is a dummy variable indicating the occurrence of a crisis between horizon $t + 1$ and horizon $t + 8$. The dummy takes the value of 1 for bank equity corrections and 2 for systemic banking crises (columns 1–2); 1 for either a systemic banking crisis or a large bank equity correction (columns 1–2). Debt-to-GDP ratio is expressed in percentage deviation from an HP-filtered trend. Significance levels at 10, 5, and 1 percent are denoted as *, **, and ***, respectively. R^2 refers to the pseudo R-squared. Robust standard errors are obtained using the multiway clustering (country and time) of Cameron, Gelbach, and Miller (2011).

systemic banking crisis or a large equity correction between horizon $t+1$ and horizon $t+8$: table 3, columns 3–4, presents these results. As can be seen, the debt-to-GDP channel remains significant.

Overall, the results in table 3 corroborate our initial findings.

5. Conclusion

We show that a tightening of monetary policy can have a positive impact on financial stability in the medium to long run, as it is successful in reducing leverage. However, the gains are moderate and occur approximately three years after the shock. In the short run, a tightening might instead generate financial instability, as it would cause further increases in the debt-to-GDP ratio and the probability of a financial crisis occurring. The inverse J-curve pattern indicates that monetary policy might not be sufficient to maintain financial stability in the short run and additional tools, such as macroprudential policies, may be needed.

Our results also point to the fact that initial conditions matter, as the higher the initial leverage, the larger is the increase in the likelihood of a crisis. It may therefore be dangerous to let the debt-to-GDP ratio increase substantially above trend. This in turn suggests that the current debate on the role of monetary policy in promoting financial stability might oversimplify the issue, as the benefits and costs of "leaning against the wind" depend on when the interventions starts.

In interpreting our results, we note that our analysis is based on the average response of variables over time and across countries. However, if the cross-country heterogeneity—for example, in macroprudential regulations—is not fully captured by fixed effects, the transmission mechanisms of monetary policy may differ. Similarly, structural breaks—because of, for example, financial liberalization or changes in the mandate of the monetary authority—might bias our responses.

Moreover, we are not conducting a welfare analysis to weigh the benefits and costs—improved financial stability and deviation of real activity and inflation from target, respectively—of a monetary policy tightening. In addition, we are unable to weigh short-term losses against long-term gains regarding financial stability. Still, we provide an important analysis in the debate on the cleaning versus leaning role of central banks by answering the question of whether monetary policy *could* lean against the wind and highlighting the undesired short-term effects of such policies.

References

Adrian, T., E. Moench, and H. S. Shin. 2013. "Dynamic Leverage Asset Pricing." Staff Report No. 625, Federal Reserve Bank of New York.

Aikman, D., M. T. Kiley, S. J. Lee, M. G. Palumbo, and M. N. Warusawitharana. 2015. "Mapping Heat in the U.S. Financial System." Finance and Economics Discussion Series No. 2015-059, Board of Governors of the Federal Reserve System.

Aikman, D., A. Lehnert, N. Liang, and M. Modungo. 2016. "Financial Vulnerabilities, Macroeconomic Dynamics, and Monetary Policy." Finance and Economics Discussion Series No. 2016-055, Board of Governors of the Federal Reserve System.

Ajello, A., T. Laubach, D. Lopez-Salido, and T. Nakata. 2016. "Financial Stability and Optimal Interest-Rate Policy." Finance and Economics Discussion Series No. 2016-067, Board of Governors of the Federal Reserve System.

Allen, F., and D. Gale. 2000. "Bubbles and Crises." *Economic Journal* 110 (460): 236–55.

Alpanda, S., and A. Ueberfeldt. 2016. "Should Monetary Policy Lean Against Housing Market Booms?" Working Paper No. 2016-19, Bank of Canada.

Alpanda, S., and S. Zubairy. 2017. "Addressing Household Indebtedness: Monetary, Fiscal or Macroprudential Policy?" *European Economic Review* 92: 47–73.

Angeloni, I., E. Faia, and M. Lo Duca. 2015. "Monetary Policy and Risk Taking." *Journal of Economic Dynamics and Control* 52 (March): 285–307.

Assenmacher-Wesche, K., and S. Gerlach. 2008. "Financial Structure and the Impact of Monetary Policy on Asset Prices." Working Paper No. 2008-16, Swiss National Bank.

Babecký, J., T. Havranek, J. Mateju, M. Rusnák, K. Smidkova, and B. Vasicek. 2012. "Banking, Debt and Currency Crises: Early Warning Indicators for Developed Countries." ECB Working Paper No. 1485.

Barberis, N. 2012. "Psychology and the Financial Crisis of 2007–2008." In *Financial Innovation*: *Too Much or Too Little?* ed. M. Haliassos, 15–28 (chapter 2). MIT Press.

Baron, M., and W. Xiong. 2016. "Credit Expansion and Neglected Crash Risk." *Quarterly Journal of Economics* 132 (2): 713–64.

Bebchuk, L. A., A. Cohen, and H. Spamann. 2010. "The Wages of Failure: Executive Compensation at Bear Stearns and Lehman 2000–2008." *Yale Journal on Regulation* 27 (2/3): 257–82.

Borio, C. 2014. "Monetary Policy and Financial Stability: What Role in Prevention and Recovery?" BIS Working Paper No. 440.

Bover, O., J. M. Casado, S. Costa, P. Du Caju, Y. McCarthy, E. Sierminska, P. Tzamourani, E. Villanueva, and T. Zavadil. 2016. "The Distribution of Debt Across Euro-Area Countries: The Role of Individual Characteristics, Institutions, and Credit Conditions." *International Journal of Central Banking* 12 (2): 71–128.

Cameron, A. C., J. B. Gelbach, and D. L. Miller. 2011. "Robust Inference with Multiway Clustering." *Journal of Business and Economic Statistics* 29 (2): 238–49.

Caprio, G., and D. Klingebiel. 2003. "Episodes of Systemic and Borderline Financial Crises." Mimeo, World Bank.

Chen, J., and F. Columba. 2016. "Macroprudential and Monetary Policy Interactions in a DSGE Model for Sweden." IMF Working Paper No. 16/74.

Chen, S., M. Kim, M. Otte, K. Wiseman, and A. Zdzienicka. 2015. "Private Sector Deleveraging and Growth Following Busts." IMF Working Paper No. 15/35.

Danielsson, J., H. S. Shin, and J. P. Zigrand. 2012. "Endogenous and Systemic Risk." In *Quantifying Systemic Risk*, ed. J. G. Haubrich and A. W. Lo, 73–94. University of Chicago Press and NBER. Available at http://www.nber.org/chapters/c12054.pdf.

Drehmann, M. 2013. "Total Credit as an Early Warning Indicator for Systemic Banking Crises." *Quarterly Review* (Bank for International Settlements) (June): 41–45.

Garriga, C., F. E. Kydland, and R. Sustek. 2013. "Mortgages and Monetary Policy." Working Paper No. 2013-037, Federal Reserve Bank of St. Louis.

Gelain, P., K. J. Lansing, and G. J. Natvik. 2015. "Leaning Against the Credit Cycle." Working Paper No. 4/2015, Norges Bank.

Gennaioli, N., A. Shleifer, and R. Vishny. 2012. "Neglected Risks, Financial Innovation, and Financial Fragility." *Journal of Financial Economics* 104 (3): 452–68.

———. 2013. "A Model of Shadow Banking." *Journal of Finance* 68 (4): 1331–63.

Giacomini, R., and T. Kitagawa. 2015. "Robust Inference about Partially Identified SVARs." Mimeo.

Goodhart, C., and B. Hofmann. 2008. "House Prices, Money, Credit, and the Macroeconomy." *Oxford Review of Economic Policy* 24 (1): 180–205.

Gourinchas, P. O., and M. Obstfeld. 2012. "Stories of the Twentieth Century for the Twenty-First." *American Economic Journal: Macroeconomics* 4 (1): 226–65.

Greenspan, A. 2002. "Opening Remarks." In *Rethinking Stabilization Policy.* Proceedings of a symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming, August 29–31.

Iacoviello, M. 2005. "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle." *American Economic Review* 95 (3): 739–64.

Iacoviello, M., and S. Neri. 2010. "Housing Market Spillovers: Evidence from an Estimated DSGE Model." *American Economic Journal: Macroeconomics* 2 (2): 125–64.

Jordà, O. 2014. "Assessing the Historical Role of Credit: Business Cycles, Financial Crises and the Legacy of Charles S. Peirce." *International Journal of Forecasting* 30 (3): 729–40.

Jordà, O., M. Schularick, and A. M. Taylor. 2015. "Betting the House." *Journal of International Economics* 96 (Supplement 1): S2–S18.

———. 2016. "Sovereigns versus Banks: Credit, Crises, and Consequences." *Journal of the European Economic Association* 14 (1): 45–79.

Laeven, L., and F. Valencia. 2012. "Systemic Banking Crises Database: An Update." IMF Working Paper No. 12/163.

Laséen, S., and I. Strid. 2013. "Debt Dynamics and Monetary Policy: A Note." Working Paper No. 283, Sveriges Riksbank.

Mishkin, F. S. 2011. "How Should Central Banks Respond to Asset-Price Bubbles? The 'Lean' versus 'Clean' Debate After the GFC." *Bulletin* (Reserve Bank of Australia) (June): 59–69.

Moon, H. R., and F. Schorfheide. 2012. "Bayesian and Frequentist Inference in Partially Identified Models." *Econometrica* 80 (2): 755–82.

Moon, H. R., F. Schorfheide, and E. Granziera. 2013. "Inference for VARs Identified with Sign Restrictions." NBER Working Paper No. 17140.

Pesaran, H. M., and R. Smith. 1995. "Estimating Long-Run Relationships from Dynamic Heterogeneous Panels." *Journal of Econometrics* 68 (1): 79–113.

Reinhart, C., and K. Rogoff. 2009. *This Time Is Different: Eight Centuries of Financial Folly.* Princeton University Press.

Robstad, O. 2017. "House Prices, Credit and the Effect of Monetary Policy in Norway: Evidence from Structural VAR Models." Forthcoming in *Empirical Economics*.

Schularick, M., and A. M. Taylor. 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008." *American Economic Review* 102 (2): 1029–61.

Smets, F. 2014. "Financial Stability and Monetary Policy: How Closely Interlinked?" *International Journal of Central Banking* 10 (2): 264–300.

Svensson, L. E. O. 2013. "'Leaning Against the Wind' Leads to a Higher (Not Lower) Household Debt-to-GDP Ratio." Working Paper. Available at http://www.larseosvensson.net.

———. 2014. "Inflation Targeting and 'Leaning Against the Wind'." *International Journal of Central Banking* 10 (2): 103–14.

———. 2016. "Cost-Benefit Analysis of Leaning Against the Wind: Are Costs Larger Also with Less Effective Macroprudential Policy?" NBER Working Paper No. 21902.

Taylor, J. B. 2007. "Housing and Monetary Policy." NBER Working Paper No. 13682.

———. 2010. "Getting Back on Track: Macroeconomic Policy Lessons from the Financial Crisis." *Review* (Federal Reserve Bank of St. Louis) 92 (3, May/June): 165–76.

Tsatsaronis, K., and H. Zhu. 2004. "What Drives Housing Price Dynamics: Cross-Country Evidence." *Quarterly Review* (Bank for International Settlements) (March): 65–78.

Uhlig, H. 2005. "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics* 52 (2): 381–419.

Discussion of "Monetary Policy, Private Debt, and Financial Stability Risks"

Òscar Jordà
Federal Reserve Bank of San Francisco and U.C. Davis

1. Introduction

Preventing a repeat of the global financial crisis has become a priority for policymakers. On the regulatory side, banking institutions are facing pressure to increase capital ratios, they are often confronted with restrictions on proprietary trading, and larger institutions are required to stress-test their loan portfolios, to list a few examples. On monetary policy, interest rate setting is no longer viewed as necessarily limited to achieving price stability and full employment mandates. Central banks question with renewed curiosity whether they should also respond to rapid growth in credit and asset prices by raising interest rates preemptively. The question itself is not new, and in the past the answer seemed to be "no." But the seminal paper by Schularick and Taylor (2012) and subsequent work has shown rapid growth of credit relative to GDP to be the single best harbinger of fragility in financial markets.

The paper by Bauer and Granziera (BG henceforth) should be viewed against this backdrop. It is not an attempt to determine whether or not it is optimal for central banks to add financial stability as one of the policy objectives to be managed with interest rate policy. Rather, it asks a narrower, more operational question: Would raising interest rates to reduce froth in credit and asset markets generate unwanted financial instability in the short run? The answer they provide is a surprising "yes" based on careful and sophisticated empirics.

The reason for this paradoxical result is that output initially declines more quickly than credit in response to an interest rate hike. As a result, aggregate "leverage" increases in the short run before coming down in the medium and longer run. I use "leverage" here to indicate above-trend growth in the ratio of credit to GDP in the aggregate sense, and not its more common usage in finance as the ratio of credit to assets.

Whether or not this finding is relevant for monetary policy depends on a number of factors. I will therefore structure my discussion around three ideas. First, financial crises in the BG sample are very infrequent. I will provide more context for this risk using a similar cross-section of countries observed over a longer period, but at a yearly frequency instead of quarterly. My evidence largely supports their findings. Because crises happen infrequently, the amount of “insurance” a central bank will be willing to pay will depend greatly on how severe these rare events are, a topic I will return to momentarily.

Importantly, because crises are uncommon, they are relatively difficult to predict (see Gadea-Rivas and Pérez-Quirós 2015). Preemptive policy will often be implemented to prevent crises that might not have taken place anyway. These are arguments commonly made against using monetary policy to achieve financial stability. That said, BG note that the interaction between crisis risk and above-trend credit-to-GDP growth is very non-linear. When credit is well above trend, crisis risk increases very rapidly. In such situations, the incentives to raise interest rates grow quickly—and so do the short-run risks that BG emphasize.

My second argument is to document that leverage tends to make recessions worse, financial or otherwise. This result, first reported perhaps in Jordà, Schularick, and Taylor (2013), suggests that a central bank may wish to restrict credit even if risk of financial crisis is not yet extreme. Crises may be rare but recessions are not.

The case for using monetary policy to hold back credit is stronger, it would seem. However, credit booms are also beneficial since they are associated with longer-lasting expansions. There is a welfare calculation to be made. How much more growth is there in an expansion with above-trend credit growth? How large is the economic loss from a credit boom gone bust? I cannot answer the exact welfare calculation here. But I can provide evidence comparing the output gains in the expansion with the potential losses in the recession. Viewed from this angle, it is less clear whether or not a central bank should aggressively manage credit booms through interest rate policy. As usual, BG remind us that monetary policy is about assessing trade-offs in a forward-looking manner.

Figure 1. Frequency of Financial Crises, 1870–2011

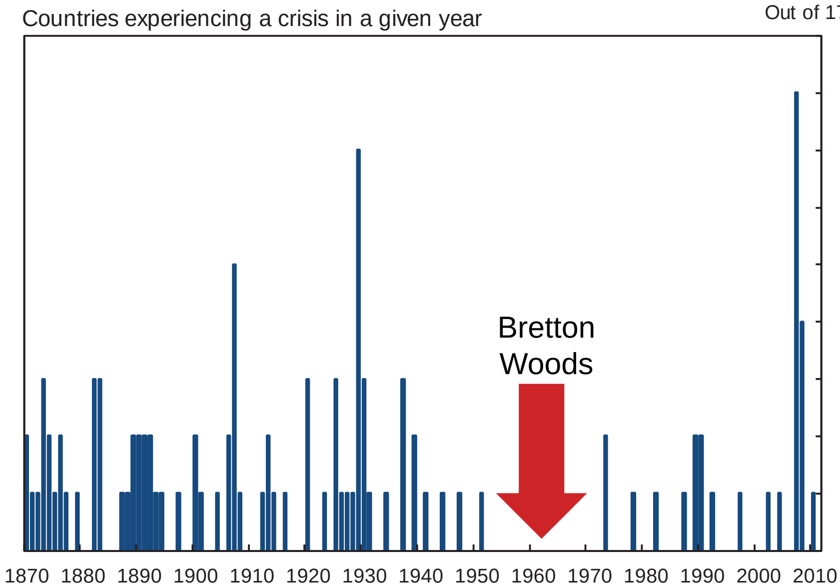

Notes: Based on seventeen advanced economies. Number of countries experiencing a crisis in a given year. Data from Jordà, Schularick, and Taylor (2017b) macrohistory database. See reference for more details.

2. Credit Booms and Financial Crisis Risk

When investigating financial crises, the BG sample, which starts in 1975:Q1, is for all intents and purposes a post–World War II sample. The reason is that there are no financial crises in advanced economies during the Bretton Woods era, from the end of World War II to 1972. Figure 1 illustrates this point by basically extending BG's figure 3 all the way back to 1870 using the Jordà, Schularick, and Taylor (2017b) database[1] (JST henceforth). Note that the data in figure 1 are aggregated into yearly frequency, and they do not include episodes of bank equity corrections, as figure 3 in BG does. This explains the minor differences between the two figures.

[1]See http://www.macrohistory.net/data/.

The key calculation in BG is that when credit to GDP grows at trend, the probability of a crisis in the following two years ranges from 1 to 12 percent. These numbers can be interpreted as an average benchmark and indeed match the unconditional probability of financial crisis in JST. Using a sample of seventeen advanced economies that nearly matches the BG country cross-section, the probability of a crisis in any country-year pair ranges from 5 percent to 2.5 percent depending on the period considered: higher in the pre–World War II period (excluding World War I), and lower in the post–Bretton Woods era.

Another way to look at these data is to classify recessions into financial or normal as Jordà, Schularick, and Taylor (2013) do. A financial crisis recession refers to a recession in which a financial crisis takes place within a two-year window of its start. Using this more narrow definition (and a slightly shorter sample including fourteen countries), full sample results suggest that about 20 percent of recessions could be associated with a financial crisis. The share declines if one excludes the post–Bretton Woods era.

Crises are rare but not entirely impossible to predict. Like Schularick and Taylor (2012), BG also find that above-trend growth of credit relative to GDP has some predictive power. Using their preferred model (column 2 in table 2 of their paper) and focusing only on systemic banking crises, the area under the receiver operating characteristic (AUROC) curve—a statistic used to assess binary event prediction that takes values between 0.5 (no better than a coin toss) and 1 (perfect classification)—is about 0.76. That is, predicting a crisis correctly 75 percent of the time will generate false positives about 25 percent of the time. These results are comparable to what Jordà, Schularick, and Taylor (2011) report using a similarly specified panel logit specification.

However, extreme deviations of credit from trend quickly raise the odds of a systemic event. Due to the non-linear nature inherent in the logit specification, BG show that deviations of credit from trend 20 percent or higher can result in probabilities of crisis increasing from the unconditional 1–12 percent range to the 30–60 percent range—a virtual certainty comparatively speaking. Accordingly, the central bank's decision to raise interest rates to prevent a crisis depends greatly on how much above trend credit is growing relative to GDP.

Figure 2. Impulse Responses to a 1 Percent Interest Rate Hike: External Instrument Identification, 1948–2013

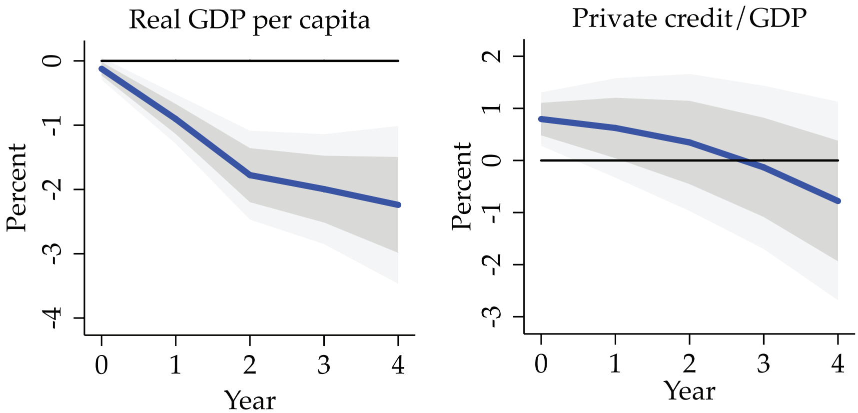

Notes: Based on Jordà, Schularick, and Taylor (2017a). The external instrument is based on the trilemma of international finance. Sample based on seventeen advanced economies. See reference for more details.

3. Credit Makes Recessions Worse, not Just Financial Crises

Results from the previous section suggest that a central bank may consider raising rates to curb credit when deviating from trend by more than some number probably north of 20 percent. Of course, a great deal more work would be necessary to find what the appropriate trend-deviation threshold is. And in practice, it is always difficult to determine where the trend is in real time. These are all relevant operational considerations.

Raising interest rates reduces credit growth, but as BG and a vast literature on monetary economics shows, it also slows economic activity down. In fact, using a completely different identification approach over a different sample than BG, Jordà, Schularick, and Taylor (2017a) find an almost identical response of output and the ratio of private credit to GDP as that reported in figure 5 of BG. Figure 2 thus provides assurance about the main result in BG.

Figure 5 in BG and figure 2 here show that in the short run, monetary policy may have the unintended consequence of *raising* crisis risk during the first year after the policy is implemented. The reason is that leverage increases over the first year, and as we saw in the previous section, more leverage increases the odds of a financial crisis. This result therefore speaks of the need for the central bank to be forward looking and to slow credit well in advance of when it reaches critical levels.

Holding back credit growth may be beneficial in other ways too. Jordà, Schularick, and Taylor (2013) show that recessions are deeper and longer lasting when credit grows above average. The sample in that paper goes back to 1870 and covers fourteen of the eighteen countries considered in BG. The main result is that recessions associated with financial crises are deeper and longer lasting.

Consider real output per capita. Normalizing relative to year 0 of the recession, in a garden-variety downturn, output declines by nearly 1.5 percent in year 1, recovers in year 2, and by year 5 is about 5 percent higher than in year 0. Contrast this to a financial crisis recession. Output declines by about 3 percent in year 2 (not year 1), and by year 5 has just about recovered its level before the contraction started. Clearly, financial crisis recessions are much worse. Figure 3 provides a visual presentation of the two recession and recovery paths just discussed.

However, regardless of type, higher rates of credit growth during the preceding expansion make matters worse. In Jordà, Schularick, and Taylor (2013), when credit to GDP grows at 3 percentage points above country-specific means on a yearly basis, output declines by over 2 percent (rather than 1.5 percent), recovers only by year 3 (rather than by year 2), and by year 5 has grown by about half of what it grows when credit hits its average growth rate. A similar picture, only slightly worse, describes what happens in financial crisis recessions.

This discussion makes evident that there may be additional advantages to holding back credit that go beyond preventing a systemic banking crisis. And such considerations would be expected to affect the "optimal" trend-deviation threshold for credit to the downside. But before concluding that central banks should be raising interest rates to guard against financial instability, it is good to remember that access to credit is one of the cornerstones to achieving higher rates of growth. The next section investigates this trade-off.

Figure 3. Cumulative Loss from Business-Cycle Peak: Normal vs. Financial Crisis Recessions, 1870–2008

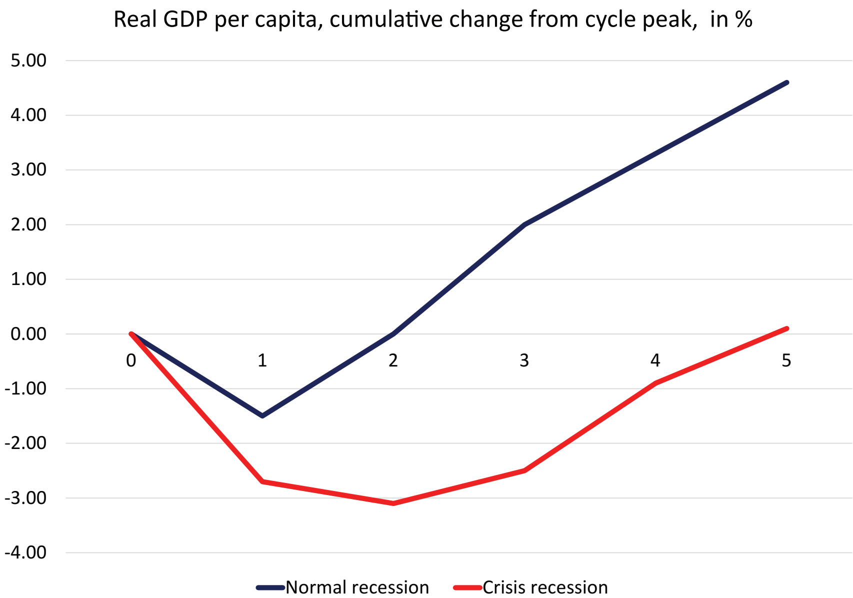

Notes: Based on seventeen advanced economies. Data from Jordà, Schularick, and Taylor (2017a). See reference for more details.

4. Credit Is an Engine of Growth

So far credit appears to have only negative consequences: it increases financial crisis risk, and it makes recessions deeper and longer lasting. Moderating credit growth plainly seems like a good idea. However, there are positive aspects to consider. After all, it is well accepted that a developed financial system is an important ingredient in achieving sustained rates of economic growth (Levine 1997).

At a business-cycle frequency, higher rates of credit growth go hand in hand with higher rates of real GDP growth per capita. Hence consider the following back-of-the-envelope calculation. Using the JST data for the post–World War II period, I break down expansions into those that experienced above-average credit growth ("high credit" expansions) versus those that were below average ("low credit" expansions). I then calculate the cumulative change in real GDP per capita experienced during the expansion as well as the loss during the subsequent recession.

High credit expansions have a cumulative growth in real GDP per capita of nearly 30 percent on average, but the subsequent recession is deeper. The cumulative real GDP per capita loss is 3.25 percent. Compare these numbers with those from a low credit expansion. These are 20 percent and 2 percent, respectively. The evidence is clear. Although recessions are made worse by high credit growth, the net gain (comparing high versus low credit expansions and subsequent recessions) is considerable—an 8.5 percent cumulative difference or about 0.7 percent faster growth per annum. That said, in a financial crisis recession, most of this advantage is lost since the recession is much deeper and lasts more than twice as long.

5. Conclusion

BG show that there is a short-run financial stability trade-off should the central bank decide to use monetary policy to moderate credit. The odds of a financial crisis increase over the year after an interest rate hike. Moreover, the odds increase dramatically when credit is already growing well above trend. The results are robust and consistent with extant results in the literature. It is unclear whether such an objective belongs in the optimal policy rule. It is clear that pursuing such an objective would require the central bank to react in a decidedly forward-looking manner. The research by BG is a welcome illustration of this important point.

References

Gadea-Rivas, M. D., and G. Pérez-Quirós. 2015. "The Failure to Predict the Great Recession—A View through the Role of Credit." *Journal of the European Economic Association* 13 (3): 534–59.

Jordà, Ò., M. Schularick, and A. M. Taylor. 2011. "Financial Crises, Credit Booms, and External Imbalances: 140 Years of Lessons." *IMF Economic Review* 59 (2): 340–78.

———. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking* 45 (s2): 3–28.

———. 2017a. "Large and State-Dependent Effects of Quasi-Random Monetary Experiments." NBER Working Paper No. 23074.

———. 2017b. "Macrofinancial History and the New Business Cycle Facts." In *NBER Macroeconomics Annual 2016*, Vol. 31, ed. M. Eichenbaum and J. Parker, 213–63 (chapter 3). Chicago: University of Chicago Press.

Levine, R. 1997. "Financial Development and Economic Growth: Views and Agenda." *Journal of Economic Literature* 35 (2): 688–726.

Schularick, M., and A. M. Taylor. 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008." *American Economic Review* 102 (2): 1029–61.

Leaning Against the Wind: Costs and Benefits, Effects on Debt, Leaning in DSGE Models, and a Framework for Comparison of Results*

Lars E.O. Svensson
Stockholm School of Economics, CEPR, and NBER

The simple and transparent framework for cost-benefit analysis of leaning against the wind (LAW) in Svensson (2017a) and its main result are summarized. The analysis of the policy-rate effects on debt in Bauer and Granziera (this issue) does not seem to contradict that the effects may be small and of either sign. The analysis of LAW in DSGE models is complicated and the results of Gerdrup et al. (this issue) may not be robust. The Svensson (2017a) framework may allow comparison and evaluation of old and new approaches and their results. As an example, it is shown that these three papers result in very different marginal costs of LAW and that a realistic policy-rate effect on unemployment is crucial.

JEL Codes: E52, E58, G01.

1. Introduction

"Leaning against the wind" (of asset prices and credit booms) (LAW for short) refers to a monetary policy that is somewhat tighter (that is, with a somewhat higher policy interest rate) than what is consistent with flexible inflation targeting without taking any effects on financial stability into account. LAW has obvious costs in terms of a weaker economy with higher unemployment and lower inflation. It has been justified by possible benefits in the form of a lower probability or smaller magnitude of future (financial) crises (Bank for

*A first version of this Commentary was presented at the Annual IJCB Conference, held at the Federal Reserve Bank of San Francisco, November 21–22, 2016. I thank the participants of the conference for comments.

International Settlements 2014, 2016; Olsen 2015; Sveriges Riksbank 2013).

In particular, the BIS has been a strong proponent of LAW, although, strikingly, without the support of a credible numerical cost-benefit analysis, as noted in an independent review of BIS research, Allen, Bean, and De Gregorio (2016, p. 19):

> Although the BIS was an early proponent of using macroprudential policies to manage the credit cycle, much of the associated BIS research and policy advice continued to focus instead on the use of LAW in monetary policy to mitigate financial stability risks. Indeed, in the outside world, this is—for better or worse—the view that the BIS is probably most associated with. Yet so far the argument for LAW seems to have cut relatively little ice with those actually responsible for setting monetary policy. In part, that is because of the lack of convincing evidence that the expected benefits outweigh the expected costs.

Instead, such cost-benefit analysis has been undertaken outside the BIS, for example, by Ajello et al. (2016), International Monetary Fund (2015), and Svensson (2014), with a rather unfavorable outcome for LAW. The extensive study in the policy paper IMF (2015) concludes:

> The question is whether monetary policy should be altered to contain financial stability risks. . . . Based on our current knowledge, and in present circumstances, the answer is generally no.

After a thorough discussion of the evidence by the Federal Open Market Committee at its April 2016 meeting, the minutes, FOMC (2016), summarize:

> Most participants judged that the benefits of using monetary policy to address threats to financial stability would typically be outweighed by the costs . . . ; some also noted that the benefits are highly uncertain.

Previously, San Francisco Federal Reserve Bank President Williams (2015) had concluded:

> Monetary policy is poorly suited for dealing with financial stability, even as a last resort.

More recently, my paper Svensson (2017a) (CB for short) (first version Svensson 2016) provides a simple and transparent framework for a cost-benefit analysis of LAW and benchmark numerical estimates of its costs and benefits. The result is that, for existing representative estimates, the costs exceed the benefits by a substantial margin. Extensive robustness tests indicate that this result is quite robust. For example, to overturn the result, the effects of LAW on the probability or magnitude of a crisis need to be more than five to forty standard errors larger than the typical empirical estimates in the literature.

2. Costs and Benefits of LAW

The framework of CB can be summarized as follows: It examines the effect of LAW on the expected discounted loss, $\mathrm{E}_1 \sum_{t=1}^{\infty} \delta^{t-1} L_t = \sum_{t=1}^{\infty} \delta^{t-1} \mathrm{E}_1 L_t$, where E_1 denotes expectations in quarter 1, δ is a discount factor, and L_t is the quarter-t loss. The expected quarter-t loss satisfies

$$\begin{aligned} \mathrm{E}_1 L_t &= (1-p_t)\mathrm{E}_1(\tilde{u}_t^{\mathrm{n}})^2 + p_t\,\mathrm{E}_1(\tilde{u}_t^{\mathrm{n}} + \Delta u_t)^2 \\ &= \mathrm{E}_1(\tilde{u}_t^{\mathrm{n}})^2 + p_t\,[\mathrm{E}_1(\Delta u_t)^2 + 2\mathrm{E}_1\tilde{u}_t^{\mathrm{n}}\,\mathrm{E}_1\Delta u_t]. \end{aligned} \tag{1}$$

In quarter $t \geq 2$ there can be either of two states of the world, either a non-crisis or a (financial) crisis, denoted n and c, respectively. There is no crisis in quarter 1. Furthermore, p_t denotes the probability of (having) a crisis in quarter t, conditional on information available in quarter 1. The variable $\tilde{u}_t^{\mathrm{n}}$ denotes the quarter-t non-crisis unemployment deviation. It is the deviation in the non-crisis state of the actual unemployment rate from the unemployment rate that is optimal under flexible inflation targeting when the possibility of a crisis is disregarded. Then the first term after the first equality sign of (1) is the probability of no crisis, $1 - p_t$, times the expected non-crisis loss, $\mathrm{E}_1 L_t^{\mathrm{n}} = \mathrm{E}_1(\tilde{u}_t^{\mathrm{n}})^2$, the expected loss if there is no crisis in quarter t.[1]

[1]As shown in CB (appendix A), the quadratic loss function $L_t = (\tilde{u}_t)^2$ is an *indirect* loss function of deviations from the optimal policy under flexible inflation targeting when the possibility of a crisis is disregarded. It takes into account a Phillips curve and the loss from inflation deviating from the inflation target.

The second term after the first equality sign is the probability of a crisis times the expected crisis loss, $\mathrm{E}_1 L_t^{\mathrm{c}} = \mathrm{E}_1(\tilde{u}_t^{\mathrm{c}})$, the expected loss if there is a crisis in quarter t. A crisis is assumed to be associated with a (possibly random) crisis increase in the unemployment rate, $\Delta u_t > 0$, so the crisis unemployment deviation is $\tilde{u}_t^{\mathrm{c}} = \tilde{u}_t^{\mathrm{n}} + \Delta u_t$, and the crisis loss is $L_t^{\mathrm{c}} = (\tilde{u}_t^{\mathrm{n}} + \Delta u_t)^2$. This crisis increase in the unemployment rate is net of any policy response during a crisis. Thus, Δu_t can be interpreted as the unemployment rate increase that is equivalent to the combination of a demand shock and any shock to the transmission mechanism of monetary policy associated with a crisis, net of the conventional and unconventional policy response at a crisis, including any restriction on the policy response such as the lower bound of the policy rate. It represents the *magnitude of a crisis*. The benchmark assumption is that it equals 5 percentage points and that the duration of a crisis is eight quarters.

The expected quarter-t loss can be rewritten as the expression after the second equality sign in (1). There, the expression in square brackets is the expected cost of a crisis, defined as the expected crisis loss less the expected non-crisis loss. The expected cost of a crisis is increasing in the expected non-crisis unemployment deviation, $\mathrm{E}_1 \tilde{u}_t^{\mathrm{n}}$.

LAW is represented by an increase in the policy rate during quarters 1–4, denoted $d\bar{\imath}_1 > 0$. This increases the non-crisis unemployment gap, $\tilde{u}_t^{\mathrm{n}}$. This in turn has two costs of LAW. The *first* is an increase in the non-crisis loss, $L_t^{\mathrm{n}} = (\tilde{u}_t^{\mathrm{n}})^2$. The *second*—less obvious, overlooked by the previous literature, but higher—is an increase in the crisis loss, $L_t^{\mathrm{c}} = (\tilde{u}_t^{\mathrm{n}} + \Delta u_t)^2$. A possible benefit of LAW is from a lower probability of a crisis, p_t, which will reduce the second term after the second equality in (1), the expected cost of a crisis, $p_t(\mathrm{E}_1 L_t^{\mathrm{c}} - \mathrm{E}_1 L_t^{\mathrm{n}})$. A second possible benefit is from a smaller magnitude of a crisis, Δu_t, which will reduce the expected crisis loss, $\mathrm{E}_1 L_t^{\mathrm{c}} = \mathrm{E}_1(\tilde{u}_t^{\mathrm{n}} + \Delta u_t)^2$.

For an initial zero expected non-crisis unemployment deviation ($\mathrm{E}_1 \tilde{u}_t^{\mathrm{n}} = 0$), corresponding to an initial situation of no leaning, the effect of LAW on the quarter-t expected loss is then taken to be the derivate $d\mathrm{E}_1 L_t / d\bar{\imath}_1 = \mathrm{MC}_t - \mathrm{MB}_t^p - \mathrm{MB}_t^{\Delta u_t} \equiv \mathrm{NMC}_t$, where

$$\mathrm{MC}_t = 2\, p_t \mathrm{E}_1 \Delta u_t \, \frac{d\mathrm{E}_1 u_t^{\mathrm{n}}}{d\bar{\imath}_1}, \tag{2}$$

Figure 1. The Marginal Cost, Marginal Benefits from a Lower Probability and Smaller Magnitude of a Crisis, and Net Marginal Cost

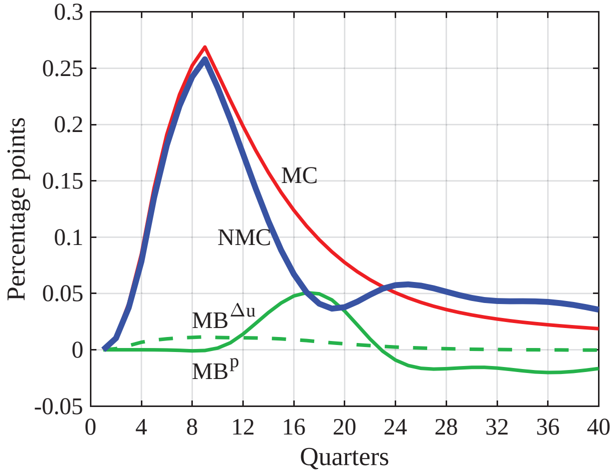

$$\mathrm{MB}_t^p = \mathrm{E}_1(\Delta u_t)^2 \left(-\frac{dp_t}{d\bar{\imath}_1}\right), \tag{3}$$

$$\mathrm{MB}_t^{\Delta u} = 2p_t\,\mathrm{E}_1\Delta u_t \left(-\frac{d\mathrm{E}_1\Delta u_t}{d\bar{\imath}_1}\right). \tag{4}$$

Here NMC_t denotes the net marginal cost of LAW, MC_t denotes the marginal cost, and MB_t^p and $\mathrm{MB}_t^{\Delta u}$ denote the marginal benefits from, respectively, a lower probability and a smaller magnitude of a crisis. Furthermore, $d\mathrm{E}_1u_t^{\mathrm{n}}/d\bar{\imath}_1$, $dp_t/d\bar{\imath}_1$, and $d\mathrm{E}_1\Delta u_t/d\bar{\imath}_1$ denote the policy-rate effects on (that is, the impulse responses of), respectively, the non-crisis unemployment rate, the probability of a crisis, and the magnitude of a crisis. In particular, the benchmark estimate of the policy-rate effect on the non-crisis unemployment rate is a representative hump-shaped impulse response that reaches a maximum of about 0.5 percentage point in quarter 6 and then slowly falls back to the baseline.

The main result of CB is illustrated in figures 1 and 2 (CB, figures 5 and 6). Figure 1 shows the marginal cost, marginal benefits, and net marginal cost of a 1 percentage point higher policy rate during quarters 1–4 than what is optimal under flexible inflation targeting when the possibility of a financial crisis is disregarded. The marginal cost, (2), is positive and substantial because of the second cost of LAW, that the expected crisis loss for a given magnitude of a crisis, $\mathrm{E}_1 L_t^c = \mathrm{E}_1(\tilde{u}_t^n + \Delta u_t)^2$, is higher with LAW. It is higher because the crisis would occur when the economy is initially weak, with higher unemployment and lower inflation, due to LAW. For the first cost of LAW, the marginal cost is zero when the initial non-crisis unemployment deviation is zero.[2]

The marginal benefit from a lower probability of a crisis, (3), is due to a representative estimate of a negative policy-rate effect on real debt growth and the effect estimated by Schularick and Taylor (2012) of real debt growth on the probability of a crisis (shown in CB, figure 3). The marginal benefit peaks in quarter 17 but then falls and turns negative after quarter 23. The reason it turns negative is that monetary policy is neutral and does not affect real debt in the long run. Then lower real debt growth and a lower probability of a crisis for some years must be followed by higher real debt growth and a higher probability of a crisis.

The marginal benefit from a lower magnitude of a crisis, (4), is due to a negative policy-rate effect on the debt-to-GDP ratio and the effect estimated of debt to income by Flodén (2014) and of debt to GDP by Jordà, Schularick, and Taylor (2013) on the magnitude of a crisis (shown in CB, figure 4).[3]

Figure 2 shows the cumulative marginal cost, marginal benefits, and net marginal cost. Clearly, this marginal cost of LAW exceeds the marginal benefit by a substantial margin.

This result is robust to monetary policy being non-neutral and having a permanent effect on real debt. Then the marginal benefit from a lower probability does not turn negative as in figure 1.

[2]Svensson (2017c) discusses the (arguably less realistic) assumptions in the previous literature about the crisis loss that disregard the second cost of LAW. They imply that the marginal cost of LAW is zero in an initial situation of no leaning. This means that *some* LAW is optimal, if benefits are positive. However, it is shown that, because the marginal cost rises rather quickly, the optimal LAW is quite small and hardly economically significant.

[3]The negative policy-rate effect on the debt-to-GDP ratio used is rather large, so as to tilt the case somewhat in favor of LAW.

Figure 2. The Cumulative Marginal Cost, Marginal Benefits from a Lower Probability and Smaller Magnitude of a Crisis, and Net Marginal Cost

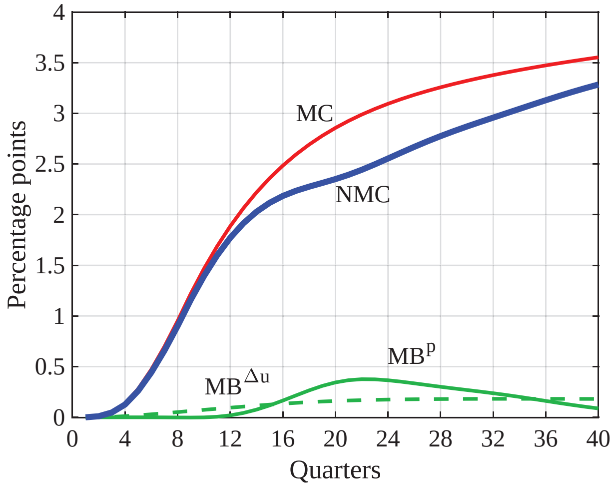

The marginal cost and benefits under a permanent effect on debt is shown in figure 3.

The result is also robust to a smaller policy-rate effect on unemployment; a credit boom with a higher probability of a crisis; a larger crisis magnitude; a longer crisis duration; less effective macroprudential policy; larger policy-rate effects on the probability and duration of a crisis; and using debt to GDP, including five-year moving averages, instead of real debt as a predictor of crises.[4] In particular, as mentioned, to overturn the result, the policy-rate effects on the probability and magnitude of a crisis need to be more than five to forty standard errors larger than the representative point estimates,

[4]Schularick and Taylor (2012) report a specification where the annual growth of real debt in the logistic function is replaced by the annual growth of debt to income. Jordà, Schularick, and Taylor (2016) use a five-year moving average of the annual growth of debt to income. Using the debt-to-GDP ratio instead of real debt or using five-year moving averages of it instead of five lags leads to a policy-rate effect on the probability of a crisis and a marginal benefit from a lower probability of a crisis similar to the one in figures 1 and 2. Thus, the cumulative marginal cost still exceeds the cumulative marginal benefit by a substantial margin.

Figure 3. The Marginal Cost, Marginal Benefits from a Lower Probability and Smaller Magnitude of a Crisis, and Net Marginal Cost; Permanent Policy-Rate Effect on Real Debt

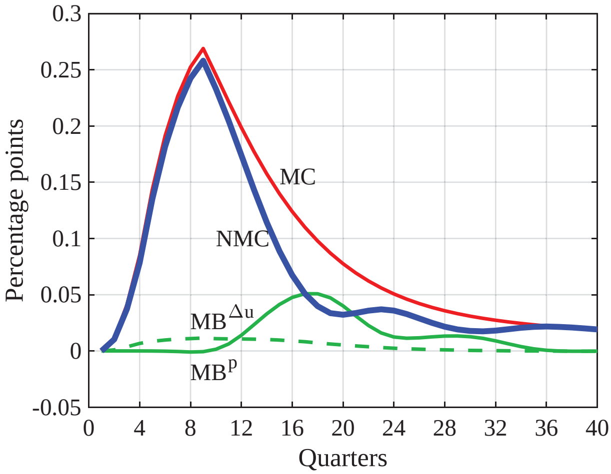

in spite of some of the estimates having large standard errors relative to the point estimates. Altogether, this indicates that the result is quite robust.

This contradicts Adrian and Liang (2018), who state that "[Svensson's] result that costs exceed benefits relies critically on assumptions about the change in unemployment in a recession or crisis, the crisis probability, and the elasticity of crisis probability with respect to the interest rate," and provide alternative assumptions that they assert would overturn the result. But Svensson (2017b) shows that their alternative assumptions are hardly realistic: they exceed existing empirical estimates by more than eleven, thirteen, and forty standard errors.

Given the simplicity and transparency of the framework and its dependence on only a few assumptions and empirical estimates of the policy-rate effects on unemployment and the probability and magnitude of a crisis, it is easy to redo the analysis with alternative assumptions or new alternative or better empirical estimates. This way the robustness of the result can be further examined.

3. The Policy-Rate Effect on Debt

As discussed in Svensson (2013), the policy-rate effect on real debt and the debt-to-GDP ratio is likely to be small and could be of either sign. The stock of nominal debt, especially mortgages with long maturities, has considerable inertia. A higher policy rate may slow down the growth of housing prices and of new mortgages, but only a fraction of the stock of mortgages is turned over each year. A higher policy rate also slows down the growth of the price level. Thus, both numerator and denominator of real debt are affected in the same direction, making the policy-rate effect on the ratio smaller and possibly of the opposite sign. This is even more the case for the debt-to-GDP ratio (a stock divided by a flow) because then not only the price level but also real GDP enter in the denominator. The policy-rate effect on the flow of nominal GDP may be larger and quicker than the effect on the stock of nominal debt. Several recent papers have indeed found empirical support for a higher policy rate *increasing* rather than decreasing the debt-to-GDP ratio (Alpanda and Zubairy 2014; Gelain, Lansing and Natvik 2015; and Robstad 2014). Nevertheless, CB uses empirical estimates according to which the policy rate has a negative effect on both real debt and the debt-to-GDP ratio, thereby stacking the cards in favor of LAW. In spite of this, the marginal cost of LAW exceeds the marginal benefit by a substantial margin.

Bauer and Granziera (this issue) (BG for short) provide an interesting and thorough examination of the policy-rate effect on real debt and the debt-to-GDP ratio. Their measure of the debt-to-GDP ratio is the percent deviation from a two-sided Hodrick-Prescott filter. They show for the point estimates of the impulse responses in their figure 5 (which are GDP-weighted averages of the point estimates for each country) that the debt-to-GDP ratio initially increases but that it later decreases. However, they report 68 percent confidence sets rather than 90 percent. If one takes a 90 percent confidence set to be 1.65 times as wide as a 68 percent, it is apparent that the effect on the debt-to-GDP ratio is not statistically significant from zero. Furthermore, their impulse responses are sensitive to the identification assumptions used as well as to alternative specifications, as shown by the robustness tests reported in their figure 7.[5] In particular, all alternatives (except using the target rate as the interest rate instead of a three-month rate) show medium-term responses of the debt-to-GDP ratio closer to zero than the benchmark response.

I find the impulse response for output in their figure 5, with a large immediate fall and a medium-term increase (which contributes to the medium-term fall in the debt-to-GDP ratio), quite unrealistic. I believe there is strong economical and empirical support (not to speak of practical policy experience) in favor of a U-shaped output response (and a hump-shaped unemployment response) from a policy tightening with a small or zero response in the current quarter.[6] Interestingly, for the added restriction of a zero initial impulse to output, the medium-term response of the debt-to-GDP ratio in figure 5 is closer to zero.[7]

As far as I can see, the results do not contradict that the policy-rate effects on real debt and the debt-to-GDP ratio are likely to be small and could be of either sign.[8]

BG also provide their own estimates of the effect of the debt-to-GDP ratio. By combining this with their estimates of the policy-rate effect on the debt-to-GDP ratio, they compute the policy-rate effect

[5]The impulse responses are also apparently sensitive to using real or nominal debt instead of the debt-to-GDP ratio in the VAR. From an economic point of view, nominal debt seems to be the natural variable, given that debt is specified and displays inertia in nominal terms. See also Brunnermeier et al. (2017).

[6]The representative output impulse response underlying figures 1 and 2 has output (translated from the unemployment response with an Okun coefficient of 2) being zero in quarter 1, falling to about 1 percent below the baseline in quarter 6, and then slowly rising back to the baseline.

[7]Looking at the impulse response of the interest rate in figure 5, with the interest rate turning negative in the medium term and the inflation response being small, one might think that the implied response of the real interest rate is negative in the medium term and that the policy shock is actually expansionary in the medium term. However, the inflation is apparently quarterly inflation at a quarterly rate. If multiplied by 4 and expressed as an annual rate and deducted from the interest rate, the resulting (ex post) quarterly real rate at an annual rate is actually positive and falling to close to zero in the medium term. Then the positive medium-term response still looks strange to me.

[8]Also, with a "financial cycle" as in Drehmann, Borio, and Tsatsoranis (2012) including real debt and debt to GDP, the policy-rate effect on it is likely to be small and possibly even of the opposite sign. Consistent with this, the policy-rate impact on the "leverage gap" in Juselius et al. (2016, table 3) is small, positive, and not statistically significant.

on the probability of a crisis within the next two years. Their results are shown in their figure 9, for the debt-to-GDP ratio being initially on trend as well as, respectively, 5 and 10 percent above trend. For a systemic banking crisis, shown in panel A, consider the case for the debt-to-GDP ratio initially on trend. Then the probability of a crisis (within the next two years) is initially 6 percent. This can then be taken to be the probability of a crisis without LAW. With LAW, the probability of a crisis first rises above 6 percent to a maximum at 7.8 percent in quarter 5, then falls back to about 6 percent in quarter 9, and finally falls below 6 percent after quarter 9. (Actually, the probability of a crisis eventually falls to close to zero, which is clearly implausible and raises some doubts about the empirical results.)

Thus, the point estimate of the probability of a crisis is higher under LAW for several quarters. More importantly, the 68 percent confidence sets in figure 9 are large and growing with the horizon. The arguably more relevant 90 percent confidence sets are even larger. All movements are well within the confidence sets, and any shift is not statistically different from zero. This is consistent with the observation by some FOMC (2016) participants that "the benefits are highly uncertain."

As far as I can see, these results do not provide support for benefits from LAW that might match or dominate the costs shown in figures 1 and 2.

4. LAW in a DSGE Model

An alternative approach to examine LAW is to introduce LAW in a DSGE model. It is of course of some interest to examine the consequences of a systematic policy of LAW, including how it might affect private-sector expectations. But such examination of a systematic policy requires a complicated DSGE model with, in practice, many assumptions and complicated parameter estimation. In particular, the mechanisms and channels through which monetary policy might affect the probability and magnitude of a crisis should ideally be included in the model, increasing the complexity of the model and of the corresponding estimation and calibration. The results will be heavily model dependent and not very robust. It is always possible to construct a model in which some LAW is optimal; the crucial and difficult questions are how reasonable and realistic the model is and how robust the results are.

Gerdrup et al. (this issue) (GHKM for short) study the consequences of LAW in a DSGE model of Justiniano and Preston (2010), a model estimated on data from Australia, Canada, and New Zealand. The role of credit is not integrated in the model, but credit developments are modeled as a block separate from the core model. Credit only affects the probability and magnitude of a crisis and has no other effects. Importantly, the private sector does not take the possibility of a crisis into account. The effects of five-year real credit growth on the probability and magnitude of a crisis are calibrated based on a sample of OECD countries. The policy-rate effect on credit is calibrated to match the response to a monetary policy shock in the structural VAR model for Norway in Robstad (2014). A (financial) crisis is represented by a negative demand shock.

Monetary policy is assumed to follow a Taylor-type policy rule. LAW is represented by the policy rate responding to pre-crisis real credit growth when it is positive but not when it is negative. Thus, LAW is asymmetric. The authors choose the response coefficients of the policy rule to minimize the weighted sum of the unconditional variances of inflation, the output gap, and the first difference of the policy rate, that is, using as a loss function

$$W_0 = \mathrm{Var}[\pi_t] + \lambda_y \mathrm{Var}[y_t] + \lambda_i \mathrm{Var}[\Delta i_t], \tag{5}$$

where Var[ ] denotes the unconditional variance.

The authors report several results. In particular, for the given type of policy rule, a positive response coefficient on real credit growth is optimal. In that sense LAW is justified in the model. However, the amount of LAW is arguably quite modest, in the sense that the policy rate is on average only 18 basis points higher with LAW than without.

As noted above, figures 1 and 2 show marginal costs of LAW that exceed the marginal benefits by a large margin. It would be very interesting to see what the these figures look like for the authors' model, in order to better understand the authors' results, but unfortunately all information required is not provided.

For several reasons, I believe the robustness of the authors' results can be doubted. First, as noted, the results of a DSGE model depend on a usually long list of specification choices and estimated and calibrated parameters. The more complex the model, the more of a black box it is, and the less clear it is how sensitive it is to alternative specifications and parameters. The Justiniano and Preston (2010) model has no less than forty-seven estimated and four calibrated parameters. Furthermore, the data behind the calibration come from several rather different countries, raising the issue of consistency among parameters. Given this, assessing the robustness of the results is not easy.

Second, more specifically, one obvious test of a model is how reasonable and realistic its impulse responses are. The paper's figure 3 shows that this model has the property that the policy rate has its largest effects on current output and inflation. Such impulse responses are in my mind highly unrealistic. Using realistic impulse responses is crucial to getting realistic estimates of the two costs of LAW mentioned above. In particular, a strong immediate policy-rate effect on output means that, by lowering the policy rate sufficiently, the central bank can completely neutralize the output effect of the negative demand shock from a crisis. This matters for the crucial second cost of LAW. A binding lower bound for the policy rate will restrict the response, but the role of the lower bound is disregarded in the paper (for instance, the experiment with a decline in foreign interest rates has the policy rate falling by about 10 percentage points). Instead, the policy is artificially restricted by the interest rate smoothing imposed (the term $(\Delta i_t)^2$ in the loss function below) and by the restriction to a suboptimal Taylor rule. Without these restrictions, in this model "cleaning" seems to be an effective way to manage crises. Given this, it would be more relevant and interesting to see what optimal policy in the model looks like, with and without a lower bound for the policy rate.[9]

[9] There is a general problem with Taylor-type rules that is overlooked by many papers examining monetary policy with such rules. Taylor-type rules are normally suboptimal; they have too few arguments. Optimal policy requires responding to all relevant state variables, including shocks. Adding an argument to a suboptimal policy rule, as is common in many papers on monetary policy, means that the set of arguments better span the space of state variables and shocks. It is therefore not surprising if adding an argument, normally any argument correlated with left-out state variables, reduces the loss.

Third, there is reason to doubt that the authors use the right loss function. The authors start from the loss function

$$W_0 = \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda_y y_t^2 + \lambda_i (\Delta i_t)^2]$$

and state that when β approaches unity the loss function approaches the weighted sum of unconditional variances in (5). However, the result of taking the limit is

$$\begin{aligned} W_0 &= \lim_{\beta \to 1-} \mathrm{E}_0 \sum_{t=0}^{\infty} (1-\beta)\beta^t [\pi_t^2 + \lambda_y y_t^2 + \lambda_i (\Delta i_t)^2] \\ &= \mathrm{E}[\pi_t]^2 + \lambda_y \mathrm{E}[y_t]^2 + \lambda_i \mathrm{E}[\Delta i_t]^2 + \mathrm{Var}[\pi_t] \\ &\quad + \lambda_y \mathrm{Var}[y_t] + \lambda_i \mathrm{Var}[\Delta i_t], \end{aligned} \tag{6}$$

where E[ ] denotes the unconditional mean. In order to use (5) one first has to show that the unconditional means $\mathrm{E}[\pi_t]$ and $\mathrm{E}[y_t]$ in (6) are zero ($\mathrm{E}[\Delta i_t]$ is trivially zero). With an asymmetric policy rule and a private sector that disregard the possibility of a crisis, the unconditional means may not be zero. Indeed, the previous Norges Bank working paper version, Gerdrup et al. (2016, table A.5) reports that the unconditional means of inflation and the output gap are lower for an asymmetric LAW than for a symmetric LAW. Thus, the loss from the unconditional means may need to be considered.

In general, an asymmetric LAW policy with a higher policy rate implies an equilibrium with lower average inflation and a lower average policy rate. To see this, take the simplest possible LAW policy,

$$i_t = r + \pi_t + \theta(\pi_t - \pi^*) + \delta, \tag{7}$$

where r denotes the average real interest rate, π^* denotes a fixed inflation target, and $\delta > 0$ denotes a constant increase in the policy rate. Take the unconditional mean of (7),

$$\mathrm{E}[i_t] = r + \mathrm{E}[\pi_t] + \theta(\mathrm{E}[\pi_t] - \pi^*) + \delta. \tag{8}$$

Assume that the Fisher equation holds, $\mathrm{E}[i_t] = r + \mathrm{E}[\pi_t]$; average inflation expectations are assumed to equal average inflation. Using this in (8) implies

$$\mathrm{E}[\pi_t] = \pi^* - \frac{\delta}{\theta} < \pi^*,$$

$$\mathrm{E}[i_t] = r + \pi^* - \frac{\delta}{\theta} < r + \pi^*.$$

For $\delta > 0$, average inflation is lower than the inflation target. Furthermore, the average policy rate is correspondingly lower, not higher. Everything else equal, this makes the economy more vulnerable to a binding lower bound for the policy rate. Indeed, if LAW with a bias towards a higher policy rate is understood by the private sector and incorporated into expectations, it is like having a lower inflation target.

Fourth, the authors report that, if the crisis magnitude is exogenous, the (restricted) optimal policy is to lean *with* the wind. Thus, the policy-rate effect on the magnitude of a crisis is apparently larger and more important than the effect on the probability. This stands in stark contrast to the result in CB, where, also for the case of monetary non-neutrality and permanent debt in figure 3, the marginal benefit from a smaller magnitude of a crisis, (4), is significantly smaller than the marginal benefit from a lower probability, (3). In particular, the small policy-rate effect on the magnitude follows rather directly from the estimates of Flodén (2014) and Jordà, Schularick, and Taylor (2013), even when a rather large estimate of the policy-rate effect on the debt-to-GDP ratio is used (so as to tilt the case in favor of LAW) (CB, figure 4). The authors do not provide any explanation for their result.

Fifth, the authors assume that five-year real credit growth is the variable affecting the probability and magnitude of a crisis and they calibrate the crucial effect on the magnitude to that in Jordà, Schularick, and Taylor (2013). But the latter use credit-to-GDP growth instead of real credit growth. Importantly, the authors do not demonstrate that their results are robust to replacing real credit with credit to GDP.[10] Furthermore, the results might be sensitive to what credit variable enters the Taylor-type rule, but no corresponding robustness tests are provided. The paper now has quarterly real credit growth in the policy rule. Given that five-year real credit growth is the variable that is assumed to affect the probability of a crisis, one might think that the policy rate should respond to that variable. But the previous working paper version, Gerdrup et al. (2016,

[10] CB (section 4.9) reports that its result are robust to using debt to GDP and five-year moving averages instead of real debt.

footnote 2), reports that the policy rate responding to five-year real credit growth leads to equilibrium indeterminacy. Gelain, Lansing, and Natvik (2015) also show that responding to the debt-to-GDP ratio or the real debt level leads to indeterminacy. There is thus reason to doubt the robustness of the policy rule used. Deriving the unrestricted optimal policy might perhaps have thrown some light on these issues.

Sixth, suppose that the probability and magnitude of a crisis depend on the effect of real debt growth, as in the paper. Suppose that monetary policy is neutral and has no effect on real debt in the long run. Then, if monetary policy manages to reduce real debt growth for some period, after that period real debt growth must eventually increase to reach the same level of real debt in the long run. Then the probability and magnitude of a crisis is shifted between periods, but the effect on the average probability and magnitude is small or zero.[11] If instead monetary policy is neutral and there is a permanent effect on real debt, the average probability and magnitude may be affected.

Indeed, the impulse response of five-year real debt growth in the paper's figure 3 shows real debt growth falling to almost -0.8 percent during quarters 8–20 and then rising back to zero, but not becoming positive. That is, there is implicit monetary non-neutrality and a permanent effect on real debt. The authors do not discuss this issue and provide no rationale for monetary non-neutrality. Any reasonable microfoundations for debt would normally imply monetary neutrality.[12]

Seventh, the paper assumes that only the central bank but not the private sector takes the possibility of a crisis into account. The previous working paper version of the paper (section 7) notes that the case for LAW is weaker if the private sector internalizes the possibility of a crisis. Given this, an information campaign, including a regular financial stability report, might increase the private sector's awareness of a crisis and be a substitute for LAW. Indeed, given the current discussion of financial stability issues and risk in the media

[11] If the logistic function for the probability of a crisis is sufficiently non-linear and increases and decreases in the probability are finely timed, the net effect on the average probability need not be zero, but it is still likely to be small.

[12] The same consequences of monetary neutrality and non-neutrality apply if the probability and magnitude depend on the growth of the debt-to-GDP ratio.

and among policymakers, it is not obvious that the private sector would underestimate the probability of a crisis.

Eighth, all credit in the model is by assumption "bad credit" in the sense of increasing the probability and magnitude of a crisis and having no beneficial effects on the economy. A point with a DSGE model is normally to provide better microfoundations, but there are no microfoundations for credit in this model. If credit is incorporated into the model and provides some benefits, the costs of LAW would increase. For example, Alpanda and Ueberfeldt (2016) incorporate credit in their model and show that costs of LAW, taking into account that it hurts borrowers, normally exceed the benefits from a lower probability of crises.

Furthermore, when the role of credit is integrated with microfoundations in the model, the relative effectiveness of monetary policy and macroprudential policy can be assessed. For example, Korinek and Simsek (2016) show that macroprudential policies can be quite effective in dealing with excess household debt and that interest rate policies are likely to be inferior in this respect. Alpanda and Zubairy (2014) incorporate housing and household debt in a DSGE model and find that macroprudential measures to reduce household debt are more effective and less costly than monetary policy. Chen and Columba (2016) analyze the effects of macroprudential and monetary policies and their interactions in an estimated DSGE model tailored to Sweden. They find that demand-side macroprudential measures are more effective in curbing household debt ratios than monetary policy, and they are less costly in terms of foregone consumption. Indeed, when credit is introduced as in Gerdrup et al. (this issue), it is trivial that a macroprudential policy restricting credit is preferable to monetary policy (see the Svensson 2012 comments on Woodford 2012).

In summary, for several reasons mentioned, the authors' results are unlikely to be robust.

5. A Framework for Comparison of Results

The simplicity and transparency of the CB framework—in particular, its dependence on only a few assumptions and estimates—make it a possible candidate for a comparison of different results about the costs and benefits of LAW. From (2)–(4) it follows that, for a comparison, what is required are assumptions about or estimates of (i) the probability of a crisis in future quarters, (ii) the expected magnitude of a crisis, (iii) the expected duration of a crisis, and the policy-rate effects on (iv) unemployment or output, (v) the probability of a crisis, and (vi) the magnitude of a crisis. With these building blocks, the marginal costs of LAW and the marginal benefits from a lower probability and smaller magnitude of a crisis can be constructed and compared and robustness and sensitivity analysis conducted.

Also for approaches that examine LAW in DSGE models, it seems that it should always be possible to construct these building blocks, including impulse responses of unemployment or output and of the probability and magnitude of crises. A comparison within this framework should at least provide some information that helps in understanding and evaluating different approaches and results.

An extensive analysis of different new and old approaches and results with the use of the CB framework is left for a separate paper. Here I will just show an example for the marginal cost of LAW. Above I noted that I found the impulse responses of output to the policy rate in BG and GHKM to be highly unrealistic. Here we will see what marginal costs of LAW they imply and how these differ from the CB benchmark marginal cost.

In figure 4, the hump-shaped solid black line shows the CB benchmark policy-rate effect on (that is, the impulse response of) the unemployment rate ($d\mathrm{E}_1 u_t^{\mathrm{n}}/d\bar{\imath}_{it}$).[13] The solid green line rising from zero to above 6 percent in quarter 9 and then falling back to 6 percent shows the CB benchmark probability of (having) a crisis for each quarter (p_t). It is given by a Markov process but is approximately equal to the sum of a crisis start in the last eight quarters, under the assumptions that there is no crisis is quarter 1, that the benchmark quarterly probability of a crisis start is 0.8 percent, and that the duration of crises is eight quarters.

From (2) and the benchmark assumption that the crisis increase in the unemployment rate (Δu) is 5 percent it follows that the marginal cost is simply ten times the product of the probability of a crisis

[13]Colors appear in the online version, but the lines are also labeled and easily identified in printed versions.

Figure 4. The Probability of a Crisis, the Policy-Rate Effect on the Non-crisis Unemployment Rate, and the Marginal Cost of LAW for Svensson (2017a) (Solid Lines) and Gerdrup et al. (2017) (Dashed Lines)

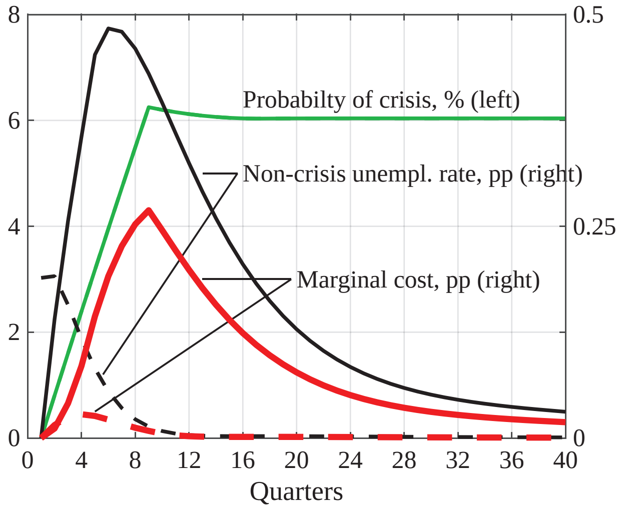

and the policy-rate effect on the unemployment rate. The hump-shaped solid red line then shows the CB benchmark marginal cost, thus calculated as ten times the product of the solid green and black lines. It is the same marginal cost as in figures 1 and 3. The area under the marginal cost is the cumulative (undiscounted) marginal cost of LAW, shown in figure 2.

The dashed black line in figure 4 shows the inferred GHKM policy-rate effect on the unemployment rate. It is inferred from the output-gap impulse response to a monetary policy shock in their figure 3 by a change of sign and division by an Okun coefficient of 2.[14] The CB benchmark probability of a crisis is assumed to apply also to GHKM. It is approximately consistent with their assumption of an annual steady-state probability of a crisis start of 3.3 percent

[14]The policy-rate impulses in CB and GHKM are somewhat different, but a correction for this would not change the conclusions.

Figure 5. The Probability of a Crisis, the Policy-Rate Effect on the Non-crisis Unemployment Rate, and the Marginal Cost of LAW for Svensson (2017a) (Solid Lines) and Bauer and Granziera (2017) (Dashed Lines)

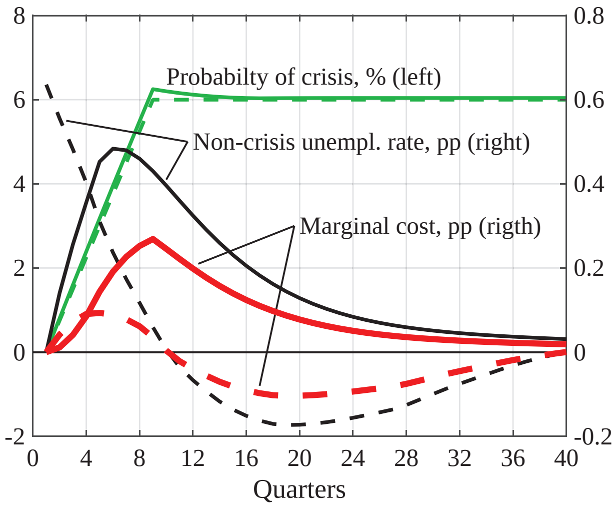

and an average crisis duration of two years. The CB benchmark crisis magnitude of 5 percentage points is also used. It is approximately consistent with the unemployment increase in their figure 1.

The dashed red line shows the resulting GHKM marginal cost, calculated as ten times the product of the solid green and dashed black lines. It is dramatically different from the CB one. The fact that the policy-rate effect on the unemployment rate is quickly falling to zero while the probability of a crisis rises from zero results in the product being very small. Whereas the CB cumulative marginal cost up to quarter 40 is 3.6, the GHKM one is only 0.19. Clearly, whether or not the policy-rate effect on unemployment is hump shaped has a big impact on the marginal cost of LAW.

In figure 5, the dashed black line shows the inferred BG policy-rate effect on the unemployment rate, inferred by a change of sign and division by an Okun coefficient of 2 of the output impulse response in their figure 5. The dashed green line shows a linear approximation of the probability of a crisis, consistent with their assumption of a steady-state probability of a crisis within eight quarters of 6 percent and thus a quarterly probability of a crisis start given by $6/8 = 0.75$ percent.

The dashed red line shows the resulting BG marginal cost. In particular, the negative policy-rate effect on unemployment after quarter 10 implies that the marginal cost is negative after quarter 10. Indeed, the cumulative marginal cost up to quarter 40 is negative, –1.3, not positive. The cumulative marginal cost is actually a cumulative marginal benefit. The importance for the marginal cost of a realistic and empirically supported policy-rate effect on the non-crisis unemployment rate is clear.

6. Conclusions

The debate about costs and benefits of LAW is likely to continue, in spite of considerable evidence that the costs of LAW exceed the benefits by a substantial margin. New papers with different approaches and theoretical and empirical results from the old ones are likely to be written. It is advantageous if the different old and new approaches and estimates can be compared and understood in a simple and transparent framework. With only a few building blocks required, the simple and transparent framework of Svensson (2017a) is a possible candidate for such comparison and understanding of different results about the costs and benefits of LAW.

References

Adrian, T., and N. Liang. 2018. "Monetary Policy, Financial Conditions, and Financial Stability." Forthcoming in *International Journal of Central Banking.*

Ajello, A., T. Laubach, D. Lopez-Salido, and T. Nakata. 2016. "Financial Stability and Optimal Interest-Rate Policy." Finance and Economics Discussion Series No. 2016-067, Board of Governors of the Federal Reserve System.

Allen, F., C. Bean, and J. De Gregorio. 2016. "Independent Review of BIS Research: Final Report." Report. Available at www.bis.org.

Alpanda, S., and A. Ueberfeldt. 2016. “Should Monetary Policy Lean Against Housing Market Booms?” Working Paper No. 2016-19, Bank of Canada.

Alpanda, S., and S. Zubairy. 2014. “Addressing Household Indebtedness: Monetary, Fiscal or Macroprudential Policy?” Working Paper No. 2014-58, Bank of Canada.

Bank for International Settlements. 2014. *84th Annual Report.* Basel: Bank for International Settlements.

———. 2016. *86th Annual Report.* Basel: Bank for International Settlements.

Brunnermeier, M., D. Palia, K. A. Sastry, and C. A. Sims. 2017. “Feedbacks: Financial Markets and Economic Activity.” Working Paper.

Chen, J., and F. Columba. 2016. “Macroprudential and Monetary Policy Interactions in a DSGE Model for Sweden.” IMF Working Paper No. 16/74.

Drehmann, M., C. Borio, and K. Tsatsoranis. 2012. “Characterising the Financial Cycle: Don’t Lose Sight of the Medium Term!” BIS Working Paper No. 380.

Federal Open Market Committee. 2016. “Minutes of the Federal Open Market Committee, April 26–27.” Technical Report, Board of Governors of the Federal Reserve System.

Flodén, M. 2014. “Did Household Debt Matter in the Great Recession?” Supplement to blog post on Ekonomistas.se. Available at www.martinfloden.net/files/hhdebt_supplement_2014.pdf.

Gelain, P., K. J. Lansing, and G. J. Natvik. 2015. “Leaning Against the Credit Cycle.” Working Paper No. 4/2015, Norges Bank.

Gerdrup, K. R., F. Hansen, T. Krogh, and J. Maih. 2016. “Leaning Against the Wind when Credit Bites Back.” Working Paper No. 9/2016, Norges Bank.

International Monetary Fund. 2015. “Monetary Policy and Financial Stability.” Staff Report.

Jordà, Ò., M. Schularick, and A. M. Taylor. 2013. “When Credit Bites Back.” *Journal of Money, Credit and Banking* 45 (s2): 3–28.

———. 2016. “The Great Mortgaging: Housing Finance, Crises and Business Cycles.” *Economic Policy* 31 (85): 107–52.

Juselius, M., C. Borio, P. Disyatat, and M. Drehmann. 2016. "Monetary Policy, the Financial Cycle and Ultra-Low Interest Rates." BIS Working Paper No. 569.

Justiniano, A., and B. Preston. 2010. "Can Structural Small Open Economy Models Account for the Influence of Foreign Disturbances?" *Journal of International Economics* 81 (1): 61–74.

Korinek, A., and A. Simsek. 2016. "Liquidity Trap and Excessive Leverage." *American Economic Review* 106 (3): 699–738.

Olsen, Ò. 2015. "Integrating Financial Stability and Monetary Policy Analysis." Speech at the London School of Economics, London, April 27.

Robstad, Ò. 2014. "House Prices, Credit and the Effect of Monetary Policy in Norway: Evidence from Structural VAR Models." Working Paper No. 5/2014, Norges Bank.

Schularick, M., and A. M. Taylor. 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008." *American Economic Review* 102 (2): 1029–61.

Svensson, L. E. O. 2012. "Comment on Michael Woodford, 'Inflation Targeting and Financial Stability'." *Economic Review* (Sveriges Riksbank) 2012 (1): 33–39.

———. 2013. "'Leaning Against the Wind' Leads to a Higher (Not Lower) Household Debt-to-GDP Ratio." Working Paper, Stockholm School of Economics. Available at www.larseosvensson.se.

———. 2014. "Inflation Targeting and 'Leaning Against the Wind'." *International Journal of Central Banking* 10 (2): 103–14.

———. 2016. "Cost-Benefit Analysis of Leaning Against the Wind: Are Costs Larger Also with Less Effective Macroprudential Policy?" IMF Working Paper No. 16/3. Available at www.imf.org.

———. 2017a. "Cost-Benefit Analysis of Leaning Against the Wind." *Journal of Monetary Economics* 90 (October): 193–213. Available at www.larseosvensson.se.

———. 2017b. "How Robust Is the Result That the Cost of 'Leaning Against the Wind' Exceeds the Benefit? Response to Adrian and Liang." Working Paper (also available as CEPR Discussion Paper No. 11744), Stockholm School of Economics. Available at www.larseosvensson.se.

———. 2017c. "Leaning Against the Wind: The Role of Different Assumptions About the Costs." Working Paper, Stockholm School of Economics. Available at www.larseosvensson.se.

Sveriges Riksbank. 2013. "Financial Imbalances in the Monetary Policy Assessment." In *Monetary Policy Report July 2013*, 42–48. Stockholm: Sveriges Riksbank.

Williams, J. C. 2015. "Macroprudential Policy in a Microprudential World." Speech at the Symposium on Asian Banking and Finance, Singapore, May 28.

Woodford, M. 2012. Inflation Targeting and Financial Stability." *Economic Review* (Sveriges Riksbank) 2012 (1): 7–32.