

Assessing the Impact of Basel III: Review of Transmission Channels and Insights from Policy Models

Olivier de Bandt, Bora Durdu, Hibiki Ichiue, Yasin Mimir, Jolan Mohimont, Kalin Nikolov, Sigrid Roehrs, Jean-Guillaume Sahuc, Valerio Scalone, and Michael Straughan

The Impact of SNB Monetary Policy on the Swiss Franc and Longer-Term Interest Rates

Fabian Fink, Lukas Frei, Thomas Maag, and Tanja Zehnder

Demand or Supply? Price Adjustment Heterogeneity during the COVID-19 Pandemic

Almut Balleer, Sebastian Link, Manuel Menkhoff, and Peter Zorn

How Wages Respond to the Job-Finding and Job-to-Job Transition Rates: Evidence from New Zealand Administrative Data

Christopher Ball, Nicolas Groshenny, Özer Karagedikli, Murat Özbilgin, and Finn Robinson

Macro and Micro Consumption Drivers in the Wake of the COVID-19 Pandemic

Elisa Guglielminetti and Concetta Rondinelli

Anchored or Not: How Much Information Does 21st Century Data Contain on Inflation Dynamics?

Michael T. Kiley

Researching the Research: A Central Banking Edition

Simona Malovaná, Martin Hodula, and Zuzana Gric

Has Higher Household Indebtedness Weakened Monetary Policy Transmission?

Gaston Gelos, Federico Grinberg, Shujaat Khan, Tommaso Mancini-Griffoli, Machiko Narita, and Umang Rawat



Assessing the Impact of Basel III: Review of Transmission Channels and Insights from Policy Models <i>Olivier de Bandt, Bora Durdu, Hibiki Ichiue, Yasin Mimir, Jolan Mohimont, Kalin Nikolov, Sigrid Roehrs, Jean-Guillaume Sahuc, Valerio Scalone, and Michael Straughan</i>	1
The Impact of SNB Monetary Policy on the Swiss Franc and Longer-Term Interest Rates <i>Fabian Fink, Lukas Frei, Thomas Maag, and Tanja Zehnder</i>	53
Demand or Supply? Price Adjustment Heterogeneity during the COVID-19 Pandemic <i>Almut Balleer, Sebastian Link, Manuel Menkhoff, and Peter Zorn</i>	93
How Wages Respond to the Job-Finding and Job-to-Job Transition Rates: Evidence from New Zealand Administrative Data <i>Christopher Ball, Nicolas Groshenny, Özer Karagedikli, Murat Özbilgin, and Finn Robinson</i>	159
Macro and Micro Consumption Drivers in the Wake of the COVID-19 Pandemic <i>Elisa Guglielminetti and Concetta Rondinelli</i>	181
Anchored or Not: How Much Information Does 21st Century Data Contain on Inflation Dynamics? <i>Michael T. Kiley</i>	239
Researching the Research: A Central Banking Edition <i>Simona Malovaná, Martin Hodula, and Zuzana Gric</i>	263
Has Higher Household Indebtedness Weakened Monetary Policy Transmission? <i>Gaston Gelos, Federico Grinberg, Shujaat Khan, Tommaso Mancini-Griffoli, Machiko Narita, and Umang Rawat</i>	325

Copyright © 2024 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 bimonthly
(ISSN: 1815-4654). Online access to the publication is available free of charge
at **www.ijcb.org**.

Requests for permission to reprint material from this journal should be addressed to:

International Journal of Central Banking
Printing & Fulfillment K1-120
Federal Reserve Board
Washington, DC 20551
Phone: 202-452-3425
Fax: 202-728-5886
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

Tiff Macklem, *Bank of Canada*

Board Members

Salah Alsavaary, *Saudi Central Bank*
David E. Altig, *Federal Reserve Bank of Atlanta*
Sergio Nicoletti Altimari, *Banca d'Italia*
Roc Armenter, *Federal Reserve Bank of Philadelphia*
Kartik Athreya, *Federal Reserve Bank of Richmond*
Jan Marc Berk, *The Nederlandsche Bank*
Claudio Borio, *Bank for International Settlements*
Jan Bruha, *Czech National Bank*
Paul Castillo, *Central Reserve Bank of Peru*
Lillian Cheung, *Hong Kong Monetary Authority*
Ken Chikada, *Bank of Japan*
Laurent Clerc, *Bank of France*
Paul Conway, *Reserve Bank of New Zealand*
Jose Gabriel Cuadra Garcia, *Bank of Mexico*
Francisco G. Dakila Jr., *Bangko Sentral ng Pilipinas*
Prof. Falko Fecht, *Deutsche Bundesbank*
Marcus Fum, *Monetary Authority of Singapore*
Carlos Garriga, *Federal Reserve Bank of St. Louis*
Joseph Gruber, *Federal Reserve Bank of Kansas City*
Philipp Hartmann, *European Central Bank*
Andrew Haughwout, *Federal Reserve Bank of New York*
Mugur Isărescu, *National Bank of Romania*
Esa Jokivuolle, *Bank of Finland*
Edward Knotek, *Federal Reserve Bank of Cleveland*
Marion Kohler, *Reserve Bank of Australia*
Sharon Kozicki, *Bank of Canada*
Signe Krogstrup, *Danmarks Nationalbank*

Michael Kumhof, *Bank of England*
Ana Cristina Leal, *Bank of Portugal*
Sylvain Leduc, *Federal Reserve Bank of San Francisco*
Carlos Lenz, *Swiss National Bank*
Ye Liu, *People's Bank of China*
Andre Minella, *Central Bank of Brazil*
Gerard O'Reilly, *Central Bank of Ireland*
Eva Ortega, *Bank of Spain*
Yang Su Park, *Bank of Korea*
Michael D. Patra, *Reserve Bank of India*
Anna Paulson, *Federal Reserve Bank of Chicago*
Thórarinn G. Pétursson, *Central Bank of Iceland*
Andrea Raffo, *Federal Reserve Bank of Minneapolis*
Trevor Reeve, *Federal Reserve Board*
Sigal Ribon, *Bank of Israel*
Kasper Roszbach, *Norges Bank*
Krislert Samphantharak, *Bank of Thailand*
Çağrı Sarıkaya, *Central Bank of Turkey*
Chiara Scotti, *Federal Reserve Bank of Dallas*
Ulf Söderström, *Sveriges Riksbank*
Salah-Eddine Taleb, *Bank of Algeria*
George Tavlas, *Bank of Greece*
Dobiesław Tymoczko, *National Bank of Poland*
Hernando Vargas Herrera, *Banco de la República*
Rafael Wouters, *National Bank of Belgium*
Egon Zakrajšek, *Federal Reserve Bank of Boston*

Editorial Board

Managing Editor

Christopher J. Waller

Board of Governors of the Federal Reserve System

Co-editors

Tobias Adrian
International Monetary Fund
Ana Babus
Washington University
in St. Louis
Huberto Ennis
Federal Reserve Bank of
Richmond
Refet S. Gürkaynak
Bilkent University

Óscar Jordà
Federal Reserve Bank of
San Francisco
Keith Kuester
University of Bonn
Robin L. Lumsdaine
Kogod School of Business,
American University

Fernanda Nechio
Federal Reserve Bank of
San Francisco
Steven Ongena
University of Zurich
Enrico Sette
Banca d'Italia

Associate Editors

Diana Bonfim
Bank of Portugal
James Cloyne
University of California, Davis
Emilia Garcia-Appendini
Norges Bank
Sarah Holton
European Central Bank
Klodiana Istrefi
Banque de France

Sergio Mayordomo
Banco de España
Emmanuel Moench
Frankfurt School of Finance
Galip Kemal Ozhan
Bank of Canada
Paolo Pesenti
Federal Reserve Bank
New York
Federico Ravenna
University of Turin

Andreas Schrimpf
Bank for International
Settlements
Zeynep Senyuz
Federal Reserve Board
Nora Traum
HEC Montréal
Kenichi Ueda
University of Tokyo

Advisory Board

Franklin Allen
Imperial College London
Charles Goodhart
London School of Economics

Hyun-Song Shin
Bank for International
Settlements
John Taylor
Stanford University

Kazuo Ueda
University of Tokyo
Carl E. Walsh
University of California
Michael Woodford
Columbia University

Assessing the Impact of Basel III: Review of Transmission Channels and Insights from Policy Models*

Olivier de Bandt,^a Bora Durdu,^b Hibiki Ichiue,^c Yasin Mimir,^d
Jolan Mohimont,^e Kalin Nikolov,^f Sigrid Roehrs,^g
Jean-Guillaume Sahuc,^{a,h} Valerio Scalone,^f
and Michael Straughanⁱ

^aBanque de France

^bFederal Reserve Board

^cKeio University

^dEuropean Stability Mechanism

^eNational Bank of Belgium

^fEuropean Central Bank

^gDeutsche Bundesbank

^hUniversity Paris-Nanterre

ⁱBank of England

This paper (i) reviews the different channels of transmission of prudential policy highlighted in the literature and (ii) provides a quantitative assessment of the impact of Basel III reforms using several policy-oriented DSGE models. It shows that the long-term effects on GDP of higher capital requirements are positive when the associated benefits are accounted for in addition to their costs. However, the results crucially depend on assumptions about crisis probability and severity. For liquidity regulations, only models capturing benefits of increased liquidity (e.g., preventing bank runs) show a net benefit.

JEL Codes: E3, E44, G01, G21, G28.

*The authors would like to thank the Editor, three anonymous referees, M. Behn, M. Birn, R. Fiori, J. Schroth, and G. Sutton for their useful comments and suggestions. This work was mainly completed while Y. Mimir was working at Norges Bank. The views expressed in this paper are those of the authors and do not necessarily reflect those of the institutions to which they belong. Declarations of interest: none.

1. Introduction

In the aftermath of the 2008 global financial crisis (GFC), the Basel Committee on Banking Supervision (BCBS) developed a set of reforms—the Basel III Accord—aimed at improving regulation, supervision, and risk management within the banking sector. By requiring banks to maintain adequate liquidity ratios and keep certain levels of reserve capital on hand, Basel III addresses a number of shortcomings in the pre-crisis regulatory framework. In order to quantitatively assess the macroeconomic impact of Basel III reforms, quantitative medium-scale models have been developed, essentially by central banks and supervisory agencies that have been at the forefront in their development and application. Are the effects beneficial for the economy according to these models?

In this paper, we address this question by first reviewing the different transmission channels of prudential policy highlighted in the literature in the last 15 years and detailing how they have been introduced in dynamic stochastic general equilibrium (DSGE) models. Second, we provide insights of the effects of Basel III from policy DSGE models routinely used for policy scenarios by central banks such as the European Central Bank, the Board of Governors of the Federal Reserve System, Norges Bank, and the Banque de France. In particular, our quantitative analysis considers how a change in the regulatory environment affects the supply of loans, lending spreads, and ultimately GDP and inflation.

Such considerations were present in the initial development of Basel III, as discussed in the Basel Committee’s Long-term Economic Impact (LEI) report (BCBS 2010) and the Macroeconomic Assessment Group (MAG 2010) report, as well as BCBS (2012). However, after a decade it is useful to revisit these issues in order to take stock of the large number of developments in macroeconomic models since then, which include a much more detailed description of the interaction between the financial sector and the rest of the economy, as well as other potential trade-offs.¹ Our analysis also

¹It is worth noting, in contrast to our analysis, that the LEI and MAG rely mostly on real-sector macroeconomic models without a banking sector, and the transmission of regulation was implemented through a calibration of the transmission of higher regulatory requirements on bank lending rates (i.e., prices)

contributes to the regular assessment of the Basel III reforms by the Financial Stability Board (FSB), an international body that monitors and makes recommendations about the global financial system, and the BCBS.

We first show that a very large number of new models have been made available since BCBS (2010, 2012), highlighting various transmission channels. Policy models reflect these channels in a consistent manner, while academic models identify other channels but are not yet sufficiently operational to allow a quantitative assessment of the impact of the regulatory changes. Second, model simulations from five policy models, based on harmonized regulatory scenarios, provide novel estimates of the impacts of Basel III. Specifically, we assess the costs associated with the implementation of Basel III, and highlight potential trade-offs between short-run losses of implementing the reforms in normal times (or when crises are mild) with long-run gains, in particular when a severe crisis occurs. The variety of models and jurisdictions on which the macroeconomic impact of Basel III is assessed helps ensure the robustness of the findings. Some models do not measure the benefits, but these may be inferred by difference from the output of the models that assess both costs and benefits. Evidence on these policy models displayed in Table 1 indicates that the long-run impact of Basel III has the expected positive sign on GDP, although the effect is not large. However, this may be associated with a economic slowdown in the transition to full implementation of Basel III. Whenever the costs and benefits of regulation are both introduced into models, the effects of Basel III are generally positive on GDP.

We find that the long-run gross benefits of the Basel III framework for the euro area could be estimated between 0.6 and 1.6 percent of GDP, as measured by the difference between the net benefits found by the 3D model (Mendicino et al. 2020), 1.2 percent, or by the model by de Bandt and Chahad (2016), 0.2 percent, and the cost found by the Gerali et al. (2010) framework, -0.4 percent. Notice that these estimates are likely underestimated due to non-linear dynamics of crises that are not well captured by the underlying

assuming a full pass-through of a higher cost of capital. Since then, the academic literature has investigated the direct impact of higher requirements on loan supply (in particular, loan quantities). See Birn et al. (2020).

Table 1. Long-Run Impact of a Move from Basel II to Basel III (Solvency) in Several Policy Models

Unit	GDP % Dev.	Bank Probability of Default, % pts. Dev.	Cost of Crisis (% of GDP), % pts. Dev.
Euro Area with 3D Model (Mendicino et al. 2020)	1.2	-7.50	-2.55 ¹
Euro Area with de Bandt and Chahad (2016)	0.2	-0.29	-0.34 ¹
Euro Area with Gerali et al. (2010) Framework (Cost Approach)	-0.4	N/A	N/A
United States (Clerc et al. 2015, with U.S. Calibration)	0.9	-9.21	-3.36 ¹
Norway—Moderate Crisis Prob. and Severity ² (Kockerols, Kravik, and Mimir 2021)	-0.2	-0.16 ⁴	-0.85 ⁵
Norway—High Crisis Prob. and Severity ³ (Kockerols, Kravik, and Mimir 2021)	2.1	-1.63 ⁴	-4.39 ⁵

¹Change in bailout costs. ²Moderate sensitivity of crisis probability and severity to credit growth. ³High sensitivity of crisis probability and severity to credit growth. ⁴Change in the probability of a financial crisis. ⁵Change in the cost of a financial crisis.
Note: The move from Basel II to Basel III is measured by a 5 percentage point increase in capital requirements.

frameworks. In addition, one needs to emphasize that the results crucially depend on the assumptions regarding the magnitude and the sensitivity of bank default probability or of the financial crisis probability to credit growth. For instance, the Markov regime-switching version of the Central Bank of Norway's Norwegian Economy Model (NEMO) (Kockerols, Kravik, and Mimir 2021) indicated that, in the case of moderate sensitivity of crisis probability and severity to credit growth, Basel III implementation reduces GDP. However, when the sensitivity of crisis probability and severity to credit growth almost doubles, Basel III has positive effects on GDP. This is consistent with BCBS (2010) and Birn et al. (2020). Expectations regarding the likely impact of the regulation also play a significant role in the positive assessment of the impact of Basel III regulations.

While significant advances have been made in the modeling of solvency requirements, more research is needed assessing liquidity requirements. Most models concentrate only on the costs of liquidity, and more work is still needed to provide a full assessment of the costs and benefits—in particular, in terms of lower contagion risk. Further modeling improvements such as the inclusion of a shadow banking sector or a more systematic inclusion of non-linear effects would allow to even better assess the impact of both capital and liquidity requirements.

This paper is organized as follows. Section 2 presents the relevant economic channels for the impact of Basel III reforms that are available in the academic literature. Section 3 discusses the contributions and limitations of policy models. Section 4 discusses the results of model-based simulation exercises and compares the outputs of some of the models surveyed in the previous section and in current use by regulators. Section 5 concludes.

2. Channels of Prudential Policy

Since the GFC, macroeconomic models have been expanded to include a more comprehensive financial sector which takes into account banks' balance sheet constraints and additional transmission channels of financial shocks, incorporating the results of models developed in the banking and finance literature. These macroeconomic transmission mechanisms of financial shocks (and thereby also prudential regulation) to the real economy have been widely

discussed in the economic literature and implemented in macroeconomic models. Table 2 sets out these channels and the corresponding contributions of the literature that are summarized in this section. In addition to these channels, it should be noted that all DSGE models are forward-looking models and incorporate a strong role of expectations, whereby agents converge to the final equilibrium, at a speed depending on the existence of nominal and real rigidities and market imperfections.

Policy models or models available in the academic literature are usually built around a core element that we call “bank capital channel,” similar to a collateral channel: increase in capital requirements lead to higher prices of bank capital, leading banks to eventually reduce loan supply and thereby affect the real economy through this price externality (Section 2.1). The notion of “bank capital channel” also appears in the context of monetary policy to denote the extent to which the central bank affects lending through its impact on bank equity when there are capital requirements for banks (see also Van den Heuvel 2006). In Section 2.2 we discuss some other possible transmission channels through banks’ liquidity, differently from bank runs discussed in Section 2.3. The integration of the possibility of bank runs into economic models introduces new trade-offs for regulation, in particular liquidity regulation. Risk-taking behavior of banks (Section 2.4) or interactions with non-bank financing (Section 2.5) combined with other elements such as the bank capital channel are other important building blocks possibly changing the trade-offs of regulation. Empirical macroeconomic models have also been developed but are not discussed here (see BCBS 2021 for more details on these models).

2.1 Bank Capital Channel

Before the 2008 crisis, some models already incorporated a collateral channel for non-financial firms. Two approaches based on different agency problems were generally used. In the first approach, the costly state-verification problem (Townsend 1979) was introduced. Information asymmetry requires lenders to pay a verification cost when borrowers default. Better capitalized borrowers are less likely to default and thus pay a lower external finance premium (Carlstrom and Fuerst 1997; Bernanke, Gertler, and Gilchrist 1999).

Table 2. Selected Contributions to the Literature by Channels of Transmission

Channels	Papers
Bank Capital Channel	Kiyotaki and Moore (1997); Gerali et al. (2010); Meh and Moran (2010); Gertler and Karadi (2011); Clerc et al. (2015); Korinek and Simsek (2016); Walther (2016); Ikeda (2018); Mendicino et al. (2018, 2020); Kravik and Mimir (2019); Jeanne and Korinek (2020); Elenev, Landvoigt, and Van Nieuwerburgh (2021); Schroth (2021); Kockerols, Kravik, and Mimir (2021)
Banks' Liquidity	Covas and Driscoll (2014); De Nicolò, Gamba, and Lucchetta (2014); de Bandt and Chahad (2016); Boissay and Collard (2016); Hoerova et al. (2018); Van den Heuvel (2019); Begenau (2020); de Bandt, Lecarpentier, and Pouvelle (2021)
Bank Runs	Angeloni and Faia (2013); Gertler and Kiyotaki (2015); Miller and Sowerbutts (2018); Gertler, Kiyotaki, and Prestipino (2020); Kashyap, Tsomocos, and Vardoulakis (2020)
Risk-Taking	Martinez-Miera and Suarez (2014); Boissay, Collard, and Smets (2016); Collard et al. (2017); Martinez-Miera and Repullo (2017); Adrian and Boyarchenko (2018); Swarbrick (2019); Coimbra and Rey (2023); Adrian et al. (2020)
Interactions with Non-bank Financing	Kale and Meneghetti (2011); Fiore and Uhlig (2011, 2015); Plantin (2015); Gertler, Kiyotaki, and Prestipino (2016); Jiménez et al. (2017); Meeks et al. (2017); Crouzet (2018); Ikeda (2018); Fève, Moura, and Pierrard (2019); Martinez-Miera and Repullo (2019); Irani et al. (2021); Begenau and Landvoigt (2022); Durdu and Zhong (2023)

The second approach builds on the human capital assumption of Hart and Moore (1994). When borrowers cannot pre-commit to work and have the option to repudiate their debt, the value of capital pledged as collateral must match the value of the debt. Thus, better capitalized borrowers can attract more funds to finance their expenditures. Kiyotaki and Moore (1997) find in their seminal paper that the collateral channel has the potential to amplify business cycle fluctuations.² Since agents fail to internalize the feedback mechanism from collateral price changes, this mechanism can also be viewed as a price externality.

The bank capital channel exhibits some parallels to the collateral channel for non-financial firms. The core idea is that a certain amount of bank capital is necessary due to an agency problem between banks and their creditors (similar to the role of collateral described above), or directly due to regulatory capital requirements. Some models combine both approaches by introducing (socially costly) financial frictions based on an agency problem mitigated by bank capital requirements. In this case, regulatory bank capital requirements have both costs and benefits: they impose credit supply constraints, but mitigate risks and lower inefficiencies associated with banks' agency problem.

When the bank capital channel is active, the impact of an adverse shock to bank capital is accompanied by a drop in bank credit supply. The exact transmission mechanism depends on the type of friction considered between banks and their creditors, and on the final use of bank funding: any fall in bank capital can result directly in higher funding costs for banks or limit banks' ability to attract funds (e.g., deposits). In both cases, a decline in credit supply affects business and/or housing investment. Some models also incorporate consumer loans which directly affect consumption. The bank capital channel thus generates interactions between the financial and non-financial sectors: adverse non-financial shocks depress activity and weaken banks' balance sheets and credit supply, exacerbating the impact of the shock on overall economic activity.

²Note that the Kiyotaki and Moore (1997) paper was revisited recently by Urban (2019), who shows that the linear approximation used in the paper reportedly led to biased results and the amplification of business cycle shocks was not quantitatively important after the correct solution method was applied.

In order to extend the collateral channel to financial intermediaries, the state-verification friction was introduced between banks and their creditors, in addition to the usual friction between banks and entrepreneurs. Hence, both banks' and entrepreneurs' capital ratios drive lending spreads (Davis 2010; Hirakata, Sudo, and Ueda 2011). A double moral hazard problem between banks and their creditors, and between banks and entrepreneurs (Holmstrom and Tirole 1997) is another friction that was introduced in DSGE models (Meh and Moran 2010). Banks exert monitoring efforts, which are costly and unobserved but can incentivize entrepreneurs to behave optimally. Higher bank capital acts as a monitoring incentive for banks and determines the amount of funds banks can attract from their creditors to fund entrepreneurs' investment. A different moral hazard problem between banks and creditors arises where banks hold equity stakes in non-financial corporations and have the option to divert a fraction of their assets to these corporations (Gertler and Karadi 2011). Depositors will liquidate the bank to avoid losses where banks divert assets. Consequently, banks hold sufficient capital to assure depositors that they do not have an incentive to divert assets and attract sufficient deposits. In all these models, changes in banks' capital affect banks' credit supply. In the same vein, Clerc et al. (2015) incorporate optimizing financial intermediaries, which allocate their scarce net worth together with funds raised from saving households across two lending activities, mortgage and corporate lending. For all borrowers (households, firms, and banks), external financing takes the form of debt which is subject to default risk. Their three default (3D) model illustrates how the three interconnected net worth channels may cause financial amplification and the distortions due to deposit insurance.

The introduction of capital requirements directly into models with the bank capital channel is crucial for measuring the costs and benefits of capital regulations. Most models capture the costs of tighter capital requirements as a reduction in credit supply. In some models, regulatory capital-to-asset ratios directly affect credit volumes (Clerc et al. 2015). In other models, stricter capital requirements are transmitted through higher bank funding costs. For example, banks deviating from exogenously given capital requirements can incur a cost captured by an ad hoc penalty function (Gerali et al. 2010).

Some models also measure the benefits of capital requirements. There is an incentive for excessive risk-taking by banks where they are protected by limited liability and consumers are protected by deposit guarantees. Higher bank capital requirements reduce bank default rates and the resources lost in the liquidation process. In addition, in Clerc et al. (2015), higher default probability for banks increases the required interest rate on uninsured bank debt and raises the cost of providing loans to the real economy. When the capital ratio is too low, the probability of bank default is high. Conversely, increasing capital from a low level may lower the weighted average cost of bank funding, as the cost of uninsured bank debt decreases, implying higher steady-state bank lending and GDP. Transition effects can also be important. The costs of tighter capital requirements can be higher in the short run (as banks must reduce credit supply) while benefits emerge in the long run as banks accumulate more capital (Mendicino et al. 2020).

The bank capital channel thus opens important interaction with macroprudential policies, but also with monetary policy. The introduction of nominal rigidities and monetary policy allows the central bank to adjust interest rates downwards in the transition phase towards tighter capital requirements to mitigate any economic slowdown caused by a drop in credit supply (Mendicino et al. 2018, 2020). Unconventional monetary policy can also be modeled as central bank credit intermediation, which interacts with the bank capital channel (Gertler and Karadi 2011).

Similar price externalities can also play a role where a liquidity trap may occur at the zero lower bound (ZLB) on interest rates (Korinek and Simsek 2016). Where constrained households deleverage after a financial shock, decreases in interest rates are needed to induce unconstrained households to support aggregate demand. If the interest rate reduction is limited by the ZLB, aggregate demand is insufficient and the economy enters a liquidity trap. In this environment, households' ex ante leverage and insurance decisions are subject to such externalities.

Some three-period models with a banking sector rely on a bank capital channel that can trigger fire sales. They can provide insights into the effect of liquidity policies alongside capital regulation. Where banks' creditors have a preference for liquidity by assumption (Walther 2016) or where the amount of liquidity is not enough in the

case of a bank run (Ikeda 2018), the economy is inefficient without regulation. In these models, both capital and liquidity requirements play a role to counter these inefficiencies. In a similar framework, Jeanne and Korinek (2020) investigate how macroprudential policy should be designed when policymakers also have access to liquidity provision tools to manage crises.

2.2 *Banks' Liquidity*

The second most important channel of transmission of financial shocks to the real economy concerns banks' liquidity. Similar to changes in capital requirements, changes in liquidity requirements lead to changes in the behavior of banks affecting the real economy. In some models, a role of liquidity regulation is created by directly imposing constraints on the maturity of funding sources, or on the liquidity of assets to avoid bank runs (without explicitly modeling bank runs; see Section 2.3). Models in this part of the literature include De Nicolò, Gamba, and Lucchetta (2014) in partial equilibrium, or Covas and Driscoll (2014) or de Bandt and Chahad (2016), in general equilibrium. In the latter model, liquidity requirements are implemented through multi-period assets and liabilities. Including households with a preference for liquidity, which banks can supply through their mix of liquid and illiquid assets, can allow for an analysis of the interaction between capital and liquidity policies (Van den Heuvel 2019; Begenau 2020). Bank capital requirements directly limit the fraction of assets that can be financed with liquid deposits, while regulation requiring banks to hold more liquid assets increases the required return on bank loans and reduces credit as a consequence. These effects have different implications for the macroeconomic costs of these policies, measured as the welfare cost to households from reduced liquidity, and lost investment and production from higher costs of intermediation.

Hoerova et al. (2018) focus on the costs and benefits of liquidity regulation. The paper demonstrates the positive role that liquidity policy can have on reducing the need for lender of last resort interventions during financial crises. The authors examine the opportunity costs of liquidity policy, providing evidence for the presence of private costs to banks resulting from requirements that force these

banks to hold more liquid assets than their own preferences, which in turn has a negative impact on profitability (see also de Bandt, Lecarpentier, and Pouvelle 2021). One major finding is that the opportunity cost of liquidity regulation is small, and smaller than that of capital regulation.

Finally, the introduction of heterogeneous banks with an inter-bank market can also create a role for liquidity policy (Boissay and Collard 2016). In the interbank market, borrowing banks divert funds that lending banks cannot easily take back, creating an agency problem. As banks do not fully internalize the effect of their funding decisions, capital and liquidity regulations can address these issues. These policies can be reinforcing where capital requirements reduce risky lending decisions and liquidity requirements encourage the purchase of lower risk liquid assets (government bonds). That said, the latter effect can reduce government bond yields, increasing demand for deposits and bank leverage (and hence risk).

2.3 Bank Runs

Reducing the risk of bank runs is a major objective of prudential policies, which triggers a new channel of transmission. A theoretical approach suitable for studying crisis phenomena (including bank runs) is global coordination games of regime change (see Morris and Shin 2001). Agents take an action (e.g., withdraw deposits from a bank, or refuse to roll over short-term debt) and their incentive to act rises with the proportion of agents taking similar actions (strategic complementarity). This innovation in the global game solution methods has spurred the development of extensions to the classical bank run model by Diamond and Dybvig (1983) (see, for example, Goldstein and Pauzner 2005).

An extension to this approach includes the introduction of safe alternative opportunities for investors to provide deposits that fund bank lending in a three-period model. Investors receive a private signal about a bank's solvency and can decide to withdraw their funding. In this way, runs are predicated on expectations about a bank's solvency. Banks have to offer a risk premium to attract funding, and more liquid banks have lower premiums. Liquidity regulation

reduces the probability of a run and can be demonstrated to reduce the profits of banks to a lesser extent than for other approaches (Miller and Sowerbutts 2018). Similar approaches have been used to examine the complementarity between capital and liquidity regulation (Hoerova et al. 2018; de Bandt, Lecarpentier, and Pouvelle 2021). Another extension of the Diamond and Dybvig framework is to include endogenous funding, with banks and borrowers subject to limited liability. Banks monitor borrowers (when profitable) to ensure that they repay their loans, while depositors may choose to run based on beliefs about both banks' monitoring and resources available for those withdrawing early (Kashyap, Tsomocos, and Vardoulakis 2020).

In a different strand of the literature, authors have started investigating infinite-horizon DSGE models combining financial accelerator effects and bank runs. In this approach, the probability of a run increases with bank leverage. Expansionary shocks increase bank leverage, bank risk, and the probability of runs. A recession that constrains bank lending due to conventional financial accelerator effects raises the possibility of runs due to the associated weakening of balance sheets. When banks optimize over a finite (two-period) horizon, regulatory constraints on leverage (in particular, countercyclical capital requirements) reduce the probability of runs, stabilize the banking system, and reduce fluctuations of the economy (Angeloni and Faia 2013). Furthermore, when banks do not take into account the effect of their leverage on asset fire sales in distressed times, there is excessive leverage in (no regulation) equilibrium (Gertler and Kiyotaki 2015). Capital requirements correct this bias and reduce the probability of runs, although there is a trade-off as tighter capital requirements also reduce the level of financial intermediation. Extending this model further to a more conventional macroeconomic setting including a production sector allows for more quantitative conclusions (Gertler, Kiyotaki, and Prestipino 2020). Compared to the more conventional models discussed in Section 2.2, the models discussed above have the advantage of capturing the highly non-linear nature of bank runs in case of a financial system collapse: when bank balance sheets are strong, negative shocks do not push the financial system to the verge of collapse; when they are weak, a similar negative shock leads the economy into a crisis in which bank runs exist in equilibrium.

2.4 *Risk-Taking Channel*

The bank-capital, bank-liquidity, and bank-run channels can all account for the fact that credit is depressed during and following a financial crisis. However, empirical studies also find that credit is also often elevated prior to financial crises (Schularick and Taylor 2012), pointing to a causal link between high credit *ex ante* and the occurrence of a financial crisis *ex post*. The literature has thus also explored how risk can endogenously build up over the course of the business cycle through a bank risk-taking channel.

The bank risk-taking channel is often built on bank's limited liability and deposit insurance. Banks enjoying limited liability do not fully internalize the risks of their lending activities, and the deposit insurance implies that banks' creditors do not require a compensation for the excessive risk-taking behavior of the bank. Compared to the social optimum, banks take too much risk, and the financial sector is too fragile. Macroprudential policies can mitigate these inefficiencies by ensuring that banks internalize these risks.

Bank risk-taking can take different forms, from extending too much credit (volume of assets) to overinvesting in risky assets (quality of assets). In Section 2.1., we already described how banks' limited liability and deposit insurance affect their credit volumes and how macroprudential policies can mitigate their risk-taking incentives. In this section, we complement this discussion by focusing on the endogenous buildup of vulnerabilities and on the risk-taking channel operating through the quality of bank assets.

Martinez-Miera and Suarez (2014) develop a DSGE model where banks can invest in safe and risky assets. Risky assets are exposed to a systemic shock affecting all these assets at the same time. When the systemic risk materializes, the equity of banks investing in risky assets is destroyed, credit supply contracts, and investment falls. In their framework, risky assets have a lower expected return on average than safe assets. Investment in the risky asset is thus socially inefficient. However, they offer a higher return when the economy is in the boom phase of the business cycle. Banks thus have an incentive to invest in risky assets, as they benefit from higher returns when the systemic shock does not materialize and are protected by limited liability when it does. When making their investment decisions, banks trade off the extra yield offered by risky assets with the

potential loss of equity. In tranquil times, bank equity becomes abundant, its value declines, and banks' incentive to take risks increases. As a consequence, systemic risk-taking peaks after long duration of normal times.

Regarding the impact of regulation, higher capital ratios (i) discourage investing in the bad asset, i.e., reduce the proportion of resources going into inefficient systemic investments, and (ii) increase the demand for scarce bank capital in each state of the economy, reinforcing bankers' dynamic incentives to guarantee that their wealth (invested in bank capital) survives if a systemic shock occurs.

Collard et al. (2017) introduce a similar framework—where banks can invest in safe or risky assets—into a standard New Keynesian DSGE model to study the interaction between monetary and prudential policies. They show that macroprudential policy is effective at mitigating bank risk-taking behaviors on the quantity and quality of credit. In contrast, monetary policy only directly affects credit volumes, because the policy rate determines the cost of funding of both safe and risky assets.

A series of recent papers have also used “savings glut” (as proposed by Bernanke 2005) to model financial crisis. As a boom progresses, investors (or savers) accumulate savings, coming from higher incomes and a desire to smooth consumption. These extra savings cause a savings glut, lowering interest rates and spreads. These lower rates, in turn, induce more risky lending such that a financial crisis becomes more likely. There is additional feedback because the risks of financial instability created by a savings glut further increase the incentive to save and to bid up asset prices.

Several frictions have been used to model the relation between saving gluts and financial crisis. They explain the endogenous buildup of risk associated with a decline in interest rates and generate boom-and-bust dynamics. They show that a banking crisis can be caused by exceptionally large shocks, or by small negative shocks at the end of a boom cycle. The bust is always associated with a credit crunch, stemming for example from bank failure, interbank market freeze, or credit rationing.

In a first set of papers, the savings glut that develops during certain booms feeds bank risk-taking behavior and culminates in bank failures. Martinez-Miera and Repullo (2017) consider banks'

investment in costly monitoring of loan quality. Banks optimally choose to monitor less in boom periods since monitoring is costly and profit is lower due to lower lending spreads, just as borrowers become riskier and crises (endogenously) more likely. A similar mechanism that is also active, mentioned already in the context of the liquidity channel, is that banks relax credit standards during a boom. Coimbra and Rey (2023) assume that intermediaries differ with respect to risk tolerance (expressed as different value-at-risk constraints). When interest rates are sufficiently low, further drops in the interest rate disproportionately benefit highly leveraged intermediaries, driving out of the market the more prudent ones. Assets then concentrate in a few large and risk-tolerant institutions, which increases financial market fragility. Similarly, Adrian et al. (2020) find that lower interest rates shift conditional output risk via further risk-taking.

In a second set of papers, the intensity of the moral hazard and/or adverse problems worsen during savings gluts, potentially causing freezes in the interbank market or credit rationing for the ultimate borrower. Boissay, Collard, and Smets (2016) introduce moral hazard and asymmetric information on the interbank market with banks that differ with respect to their intermediation skills. In a frictionless interbank market, less efficient banks raise deposits and lend their proceeds to the banks that are more efficient in intermediating funds to firms. The late-cycle savings glut reduces interbank interest rates, exacerbates the moral hazard problem, and leads to a freeze of the interbank market. The interbank market freeze forces less efficient banks to lend directly to firms (instead of delegating the task to the most efficient banks) and leads to a drop in credit. Swarbrick (2019) models asymmetric information on credit risk between firms (which know the risk) and banks (which do not), leading to adverse selection and credit rationing. During a savings glut, the risk of default increases to the point where banks decide not to lend all available funds and restrict credit to safe firms.

2.5 Interactions with Non-bank Financing

While our discussion so far has focused on banks, the financial sector is a complex system where banks, non-bank financial intermediaries

(NBFI), and the bond market interact and provide different sources of external funds to borrowers. The ability to include NBFI and direct bond issuances are important modeling features that increase the relevance of macroeconomic models used in policy analysis for economies that rely comparatively less on banking finance.

A number of recent papers try to explicitly include NBFI that are unregulated (or lightly regulated) and not covered by deposit insurance. In these papers, changes in capital requirements (or shocks on regulated banks' capital) lead to substitution effects between regulated bank and NBFI loans. In a quantitative general equilibrium (real business cycle type) setting, increasing capital requirements forces regulated banks to find more expensive equity funding, leading to an expansion of the NBFI sector as it becomes relatively more profitable (Begenau and Landvoigt 2022). A further insight from this class of model is that the broader effect of higher capital requirements in the regulated banking sector on non-financial sector borrowing may be ambiguous, as reduced leverage (and capacity for lending) at regulated banks is absorbed by higher leverage (and lending) at NBFI (Durdu and Zhong 2023). In a more stylized model built on similar assumptions, tighter leverage restrictions also improve the resilience of regulated banks but lead to increased leverage (and, as a result, higher default probability) by NBFI (Ikeda 2018).

The literature also offers insights on the modeling of other NBFI characteristics. NBFI tend to finance their assets with short-term wholesale funds (e.g., Gertler, Kiyotaki, and Prestipino 2016) and have special securitization skills: they issue tradable securities backed with (bank) loans (e.g., Meeks, Nelson, and Alessandri 2017). In such setting, NBFI can be the source of financial shocks that propagate to banks and the macroeconomy. NBFI can also amplify the effect of aggregate shocks leading to a synchronized increase in economic activity and credit demand, because they help to alleviate the constraints faced by regulated banks (Fève, Moura, and Pierrard 2019).

Other papers focused on borrowers' (mostly entrepreneurs) choice between bonds and bank loans. Bond finance is cheap, but banks can mitigate information asymmetry problems through screening and monitoring and have special restructuring skills. In De Fiore and Uhlig (2011, 2015), firms can pay a cost to share

some information on their productivity with the bank. Firms with an intermediate level of productivity find it optimal to pay the information cost and opt for bank loans. The most productive firms are financed with bonds, while the least productive abstain from raising funds. When bank loans can be restructured in times of financial distress (thus avoiding inefficient liquidation of potentially viable firms), riskier firms opt for a mix between bonds and bank loans, while safer firms exclusively rely on the bond market (Crouzet 2018). In this type of model, bonds are substitutes for bank loans. Bond issuance can mitigate the impact of bank credit supply shocks. These models are, however, silent on the effect of a change in capital requirements.

It is also well-known that larger firms are more likely to have access to the bond market (for empirical evidence, see Kale and Meneghetti 2011 and the references therein). de Bandt and Chahad (2016) introduce small and large non-financial corporations when studying the impact of prudential regulations in a policy model. While small firms exclusively rely on bank loans, large firms are financed with a mix of bond and bank loans, which also generates substitution effects between bank and bank loans for the latter.

Martinez-Miera and Repullo (2019) extends this discussion to the choice between bonds, bank loans, and NBFI loans. They build a stylized model where financial intermediaries screening is costly and unobserved but mitigates entrepreneurs default risks. They find that low-risk entrepreneurs borrow from the market. Intermediate and high-risk entrepreneurs borrow from intermediaries. Their choice between bank and NBFI loans depends on the design of prudential regulations. When capital requirements are risk based, intermediate-risk entrepreneurs opt for bank loans while high-risk borrowers are financed by NBFI. Similarly to the aforementioned literature (as well as Plantin 2015), a tightening of capital requirements in the banking sector may lead to substitution effects (see empirical evidence in Jiménez et al. 2017, Irani et al. 2021). As a consequence, the optimal capital ratio may be lower compared to a situation without a shadow banking sector (or financial market). But at the same time, the overall financial system is riskier, due to the existence of an unregulated shadow banking sector. The overall macroeconomic effect of a larger shadow banking sector is still an area for additional research.

3. Contributions and Limitations of Policy Models

We now discuss how policy models used by central banks and supervisory authorities incorporate the channels presented in Section 2 and what are their limitations. In particular, we highlight similarities and differences between models developed in a representative set of jurisdictions: (i) the euro area (3D model by Clerc et al. 2015, complemented by a monetary policy channel, as described by Mendicino et al. 2020; model by de Bandt and Chahad 2016; and a modified version of Gerali et al. 2010, as described by Bennani et al. 2017), (ii) the United States (3D model by Clerc et al. 2015, with a country-specific calibration used by the Board of Governors) and (iii) Norway (Norges Bank, regime-switching version of NEMO; see Kockerols, Kravik, and Mimir 2021). We also consider models developed in other policy institutions. Darracq-Pariès et al. (2022) carry out a similar exercise but only focusing on models at the European Central Bank. We then investigate their performance in light of results highlighted in academic papers. Papers under review in the section are displayed in Table 3.

3.1 *Contributions of Policy Models*

Policy models are operational macroeconomic models used for policy simulations. They implement in a consistent fashion some of the characteristics of academic models and bring them to the data. Policy models share the following features. First, the models are estimated or calibrated to get a reasonable level of plausibility. This implies in particular that they are cast in a national accounting framework, at a quarterly frequency. They are well adapted to mirror actual economic data and trends. One should not, however, underestimate the challenges in estimating and calibrating these models in the context of setting initial steady-state banks' probability of default (BCBS 2021, Annex 1) or financial crisis probability (BCBS 2021, Annex 6). Although the default probability is endogenous in these models and the elasticity of the default probability depends on structural factors, the initial value of default probability matters. Moreover, in some of these models, financial crisis probability depends on credit growth that is endogenous, but the sensitivity

Table 3. Policy and Academic Literature on the Modeling Challenges Associated with Basel III

Challenge	A Selection of Policy Models	Other Academic References
Detailed Regulatory Policies	Gerali et al. (2010); Covas and Driscoll (2014); De Nicolò, Gamba, and Lucchetta (2014); Clerc et al. (2015); de Bandt and Chahad (2016); Mendicino et al. (2018, 2020); Kravik and Mimir (2019) (Central Bank of Norway's NEMO model); Darracq-Paries et al. (2022)	Brunnermeier and Sannikov (2014); Holden, Levine, and Swarbrick (2019); Lang and Forletta (2020); Dou et al. (2021); Elenev, Landvoigt, and Van Nieuwerburgh (2021); Schroth (2021); Jondeau and Sahuc (2022); Suarez (2022) Adrian and Boyarchenko (2018); He and Krishnamurthy (2019); Begenau (2020) Gertler, Kiyotaki, and Prestipino (2020) Badarau and Leveuge (2011); Corbae and D'Erasmo (2019); Coimbra and Rey (2023)
Non-linearity	Kockerols, Kravik, and Mimir (2021)	
Multiplicity of Channels	See line 1	
Endogenous Crisis Heterogeneous Banks	Kockerols, Kravik, and Mimir (2021)	

of crisis probability to credit growth is estimated based on single-country or cross-country data outside the model.

Second, the models include a quite high level of details, allowing a measurement of actual policy which is clearly not the case in more stylized models. For example, the 3D model includes two types of assets, and two types of liabilities, even if they are one-period debt. In the model by De Nicolò, Gamba, and Lucchetta (2014) banks may invest and issue one-period bonds, besides collecting deposits and granting loans. In the model by de Bandt and Chahad (2016), assets and liabilities are defined for more than one period, allowing a quite accurate modeling of liquidity requirements. In the Central Bank of Norway's model, NEMO, banks extend credit to both households and non-financial firms and finance themselves using household deposits and foreign debt. Household debt is long term, which reflects the long-term nature of mortgage debt.

Third, they integrate many of the channels presented in Section 2, in contrast to stylized academic models that often concentrate on a particular channel of transmission. To be more specific:

- The *bank capital channel* is present in all policy models, as solvency is a central tool of banking regulation. A key benefit of increasing capital in the 3D model—as well as for Kockerols, Kravik, and Mimir (2021) for Norway—is the expected reduction in the probability of bank failure. It also induces a reduction of bank runs for the de Bandt and Chahad (2016) model. The models reviewed by Darracq-Pariès et al. (2022) also include a strong link between bank capital and credit supply. Bank default appears a key factor affecting the response of the economy to higher capital requirements in the medium to long term. In 3D, fewer bank failures imply lower bank failure costs—both public and private. The public costs of deposit insurance and the bailing out (or resolution) of failing banks are ultimately borne (for simplicity) by all households because they are taxpayers. Other deadweight costs also affect households' consumption. The private costs are captured by the spread banks are forced to pay over the risk-free rate in order to attract debt funding. Some bank debt is uninsured and its interest rate decreases when banks are safer because debt holders no longer need to be compensated for the potential

losses. When banks are competitive, as in the 3D model, this cost reduction will be passed on to borrowers (*ceteris paribus*), stimulating economic activity. When they are facing monopolistic competition, like in NEMO, the pass-through is smaller in the short run but is ultimately passed on to borrowers. Nevertheless, when there is an occasional financial crisis, lending spreads in NEMO become higher during crisis episodes due to asymmetrically large credit supply shocks. In this case, a benefit of reducing the crisis probability (by raising capital requirements) is to have lending spreads that are lower in crisis times, at the cost of being higher in normal times.³ An effect that pushes in the opposite direction, and that is a key cost of increasing capital requirements, arises when the required rate of return on equity is higher than the cost of debt (which is usually the case and assumed in the 3D model, in NEMO, and in many other macrofinancial models). The higher return on equity means that higher capital requirements adversely affect banks' profits, which increase the spread of lending rates over deposit rates in order to achieve higher profitability and attract equity investors.

- *Liquidity and bank runs* are introduced in de Bandt and Chahad (2016), while the *liquidity channel* is also introduced in the De Nicolò, Gamba, and Lucchetta (2014) partial equilibrium model, and the Covas and Driscoll (2014) DSGE model. In the latter model, general equilibrium effects matter; notably, the increase in loan spreads following a tightening of regulation softens the initial negative effects on loan supply. But the consideration of both capital and liquidity constraints is not very frequent in the academic literature, and also not prevalent in many policy models. Impacts on bank runs are rarely introduced, except for de Bandt and Chahad (2016). For instance, for De Nicolò, Gamba, and Lucchetta (2014), liquidity requirements reduce banks' liquidity transformation

³Tighter solvency regulation in NEMO reduces the probability of occurrence of crisis periods characterized by a large increase in lending spreads, hence partially reducing the increase in the ergodic mean of lending spreads over the whole business cycle, because higher capital requirements slow down household credit growth upon which the crisis probability depends.

activity, hence lending activity, without consideration of long-run effects on financial stability.

- *Financial crisis* into policy models is introduced through a regime-switching DSGE model, as in Kockerols, Kravik, and Mimir (2021) (as implemented in Norges Bank's main macroeconomic model, NEMO).⁴ This type of policy model allows one to flexibly incorporate different types of non-linearities into large-scale quantitative macro models such as financial crises by combining typical business cycle fluctuations (normal times regime) with crises dynamics (crisis times regime). The evolution of crisis is then calibrated to replicate the dynamics of many macroeconomic and financial variables during actual crisis episodes or those under macroprudential stress-testing scenarios used at central banks. Moreover, the transition from normal times to crisis times depends on endogenous regime-switching probabilities that link credit growth (or leverage) to crisis probability. Crisis probability and severity endogenously fluctuate based on the evolution of credit growth in the model.
- The *risk-taking channel* is introduced in the 3D model through limited liability (see Section 2.4.).

Note, however, that some academic models also include multiple channels through interaction between capital and liquidity requirements. This is the case of the Adrian and Boyarchenko (2018) continuous-time model, which looks at the interaction between liquidity and solvency requirements, and concludes that liquidity requirements reduce systemic risk without impairing consumption growth. This is also the case of the Begenau (2020) model, which includes deposits in households' utility functions. Deposits are cheaper compared to equities, because they offer a convenience yield. Higher equity requirements reduce bank deposits. This is welfare decreasing for households that derive utility from deposits. However, it also decreases the deposit rate (since the marginal convenience yield increases), which lowers banks' financing costs. Higher equity

⁴The Markov regime-switching version of NEMO is solved using the Rationality In Switching Environments (RISE) toolbox developed by Junior Maih. RISE is an object-oriented MATLAB toolbox for solving and estimating non-linear regime-switching DSGE models. The toolbox is freely available for downloading at https://github.com/jmaih/RISE_toolbox.

requirements also increase banks' incentive to monitor projects, as shareholders have more "skin in the game." This lowers banks' risk and raises their average returns, leading to more credit provision. Overinvestment in low-quality projects decreases. This reduces the volatility of output and consumption and boosts their average levels. Benefits from higher and smoother consumption outweigh the costs of lower deposits. The "deposit rate" channel complements the traditional bank capital channel.

3.2 Limitations of Policy Models

Policy models face, however, some limitations.⁵ First, one criticism of operational models is that they fail to adequately capture observed economic outcomes during crisis periods. Linearized solutions to DSGE models around a steady state disregard underlying non-linear dynamics. While some models explicitly allow for the presence of occasionally binding constraints on bank capital, the simulations are performed assuming that these constraints are always binding. So by assumption there is no financial crisis in the latter models, as only the normal business cycle fluctuations are modeled. As a consequence, these models drift back towards a steady state more quickly than real economic data demonstrate. Bank capital constraints may be binding during crises but not in normal times.

The non-linear financial crisis dynamics integrated into Kockerols, Kravik, and Mimir (2021) are not fully microfounded. The evolution of crisis dynamics are determined using regime-switching structural model parameters and asymmetrically large shocks. Furthermore, in some of these models, the negative externalities that are driven by crisis probability and severity are estimated empirically outside the model. In particular, the sensitivity of crisis probability and severity to credit growth is calibrated to match the empirical estimates of crisis probability and severity from logit-type regressions and local projection models, respectively.

In contrast, several academic papers explicitly introduce non-linear dynamics as discussed in Dou et al. (2021), although they remain stylized. They analyze the full non-linear dynamics leading to uncertainty over the persistence of a crisis. In the paper by

⁵See also BCBS (2012).

Brunnermeier and Sannikov (2014) greater instability is generated once the economy moves sufficiently away from the steady state (crisis times), as equity capital is increasingly misallocated due to the uncertainty, leading to underinvestment and distorted household consumption decisions. Similarly, in He and Krishnamurthy (2019) where the amount of equity raised by an intermediary depends on uncertain equity returns subject to shocks, a negative shock when intermediaries are relatively unconstrained (“normal” times) triggers only a small decline in equity, asset prices, and investment. However, when these equity constraints are binding or likely to be binding in the near future (“crisis” times), a negative funding shock triggers a more substantial decline. In addition, Jondeau and Sahuc (2022) show that non-linearities (fire sales and bank default) in a DSGE model with two types of banks play a fundamental role in the development of a crisis. Occasionally binding constraints on bank debt (as opposed to bank capital) can generate similar non-linear dynamics and outcomes. Under normal circumstances, banks rely on debt finance whereas under financial stress borrowing constraints can bind and banks must eventually raise additional equity finance at a higher cost. This results in occasional episodes with sharp increases in spreads and deeper downturns, helping explain observed macroeconomic asymmetries such as negatively skewed aggregate investment (Holden, Levine, and Swarbrick 2020).

Recently, models including a bank capital channel with occasionally binding constraints were used to analyze the effectiveness of prudential policy. Pre-crisis capital requirements in the United States were found to be close to optimal in terms of the aggregate welfare of savers and borrowers where default and occasionally binding borrowing constraints in both the non-financial and financial sectors are present (Elenev, Landvoigt, and Van Nieuwerburgh 2021). Capital buffers are found to be more effective in restricting bank equity payouts, rather than bank lending over financial cycles, in examining the implementation of the capital conservation buffer and the countercyclical capital buffer (Schroth 2021). Although these models resolve the limitation of many DSGE models concerning the inclusion of non-linear effects, other elements and frictions are not present in these models (e.g., liquidity requirements, nominal frictions), limiting their potential usefulness for policy analysis.

It is also rather clear that, while long-run dynamics are satisfactory, modeling the nexus between real and the financial sector in these models requires further advances to capture the non-linearities associated with large adverse shocks such as financial crises. Using standard grid-based global solutions is challenging due to the curse of dimensionality. However, new techniques using insights from machine learning appear to be a promising way to handle non-linearities also in larger models. This is an important priority for future work.

Furthermore, financial crises are rare events, but when they occur they tend to be associated with heavy losses for banks. Failure to take into account non-linearities, and more particularly tail events, will thus lead to underestimation of the effects of financial crises. For instance, by studying the impact of cyclical systemic risk on future bank profitability for a large representative panel of EU banks between 2005 and 2017, Lang and Forletta (2020) show that the impact of cyclical systemic risk on the left tail of the distribution of future bank returns is an order of magnitude larger than on the median. In addition, Suarez (2022) explains that using wrong estimates of the (non-linear) causal effects of risk and policy variables on the relevant moments (mean and growth-at-risk) of the growth distribution could produce misguided macroprudential policy advice.

A second limitation is related to the level of details present in the models. Bank capital is modeled in a simple fashion. There is only one type of capital, which is accumulated only through retained earnings and, in most cases, there is no distinction between regulatory and voluntary capital buffers. In addition, most policy models are based on a representative bank (or a continuum of identical banks), while heterogeneous agent models are currently being developed, with heterogeneous agent New Keynesian (HANK) type models. Note, however, that the 3D and Gerali et al. (2010) models already include two types of households, with patient and impatient households. Patient households are savers who buy housing without a need for financing, while impatient households borrow from banks using their holdings of housing as collateral. With these features, these policy models allow for some degree of heterogeneity, even though they do not feature the rich heterogeneity in the HANK models. In addition, an interesting extension to the representative bank model is the heterogeneous bank framework, such as Badarau

and Leveuge (2011), Corbae and D’Erasmus (2019), or Coimbra and Rey (2023). A growing literature quantitatively investigates the distribution of bank balance sheet items depending on their size and portfolio composition and its interactions with solvency regulation.

Third and finally, most policy models do not include non-bank financial sector, so that, as discussed in Section 2, leakages would not be captured. Note, however, that some policy models reviewed by Darracq-Pariès et al. (2022) indicate that the existence of non-bank funding sources can help dampen the effect of the short-term tightening of loan supply in response to capital requirements.

4. Model-Based Quantitative Illustrations

This section provides the building blocks for future regulation assessments by illustrating the performance of selected policy DSGE models (e.g., the type of response they provide, distinguishing between benefits and costs) in identifying the macroeconomic impact of Basel III reforms. Specifically, we use five of the models discussed in Section 3: (i) the 3D model by Clerc et al. (2015); (ii) the 3D model complemented by a monetary policy channel, as described by Mendicino et al. (2020); (iii) the model by de Bandt and Chahad (2016); (iv) a modified version of Gerali et al. (2010), as described by Bennani et al. (2017); and (v) a regime-switching version of Norwegian Economy Model, NEMO, as described by Kockerols, Kravik, and Mimir (2021).

The calibration of these models is based on the most recent data, in order to capture the current state of the economy and to be able to perform simulations applied to the current context. In the calibration procedure, we thus assume that most of the Basel III regulatory agenda has been implemented, which is true to a large extent, and set out deep structural parameters to match the means and variances of our data in the recent period. With this calibration of deep structural parameters at hand, we can perform various types of experiments by adjusting the values of our “Basel III parameters.”

To study the macroeconomic impact of Basel III reforms, we implement several scenarios, which consider solvency and liquidity regulations. First, we increase capital requirements to capture the increase in the quantity and quality of capital requirements that the Basel III reform imposed. In a second scenario, we also consider

the fact that Basel III imposed an increase in liquidity requirements. In each case, we analyze the impact of Basel III on long-run equilibrium values of important macroeconomic variables, and study the transition to the new regime, from Basel II to Basel III.

In addition, our analysis allows us to distinguish between the costs and benefits of solvency and liquidity regulations, as shown in Tables 4 and 5. The implementation of the 3D model with different calibrations on simulations for the euro area, as well as for the United States, permits an assessment of the contribution of country/area idiosyncrasies. In contrast, the Norges Bank's model offers a different modeling perspective. All in all, Basel III appears to have the effects anticipated, in terms of positive effects on GDP and financial stability (and the exercise offers a useful quantification of these effects), although its contribution to real macroeconomic outcomes appears to be small. In particular, we find that (i) the calibration/measurement of the bank default probability (or financial crisis probability) and its evolution play a crucial role in the assessment and (ii) the expectation channel plays an important role, conditioning the transition path associated with the reforms: if economic agents anticipate that the reforms will effectively reduce the probability of bank failure or the probability of a run, this triggers, beyond the initial supply shock, a positive demand effect on GDP.

4.1 Solvency Scenario

The Basel III reforms increased banks' minimum capital ratio from 4 percent to 6.5 percent. However, there was a recognition that some hybrid forms of capital did not effectively absorb losses in the crisis. Hence Basel III also increased the share of banks' capital which comes from loss-bearing common equity tier 1 (CET1) capital. In this sense, the reforms increased not only the quantity of capital but also its quality due to its greater ability to absorb losses. The models we use do not have such a high degree of granularity. They include risk weights, but there is only one type of bank capital. However, this does not necessarily prevent us from analyzing fully the impact of Basel III regulation that includes requirements in terms of both quality and quantity of capital.

Table 4. Long-Run Impact of a 5 Percent Increase in Capital Requirements

	Expected Benefits of Regulation				Costs of Regulation	Real Macro Variables			Financial Macro Variables
	Bank Probability of Default	Cost of Crisis	Bailout Cost as % of GDP	Bank Debt Funding Cost Spread over Risk-Free Rate		Lending Spread over Bank Debt Fund. Cost	GDP	Aggregate Investment	
Unit	% Pts. Dev.	% Pts. Dev.	% Pts. Dev.	% Pts. Dev.	% Pts. Dev.	% Dev.	% Dev.	% Dev.	% Dev.
Euro Area with 3D	-7.50	NaN	-2.55	-0.59	0.34	1.19	0.29	1.45	2.55
Euro Area with de Bandt and Chahad	-0.29	NaN	-0.34	0.08	0.02	0.2	0.56	0.18	1.26
Euro Area (Cost Approach)	NaN	NaN	NaN	0.17	0.11	-0.4	-1.31	-0.45	-5.85
United States with 3D	-9.21	NaN	-3.36	-1.43	2.48	0.87	7.53	4.07	8.03
Norway (NEMO) ¹	-0.16 (*)	-0.85 (**)	NaN	NaN	0.59	-0.18	-2.96	0.57	-3.18
Norway (NEMO) ²	-1.63 (*)	-4.39 (**)	NaN	NaN	0.59	2.1	12.4	0.28	12.9

(*) Change in the probability of a financial crisis. (**) Change in the cost of a financial crisis.
Note: 3D-EA: Mendicino et al. (2020); EA-Cost Approach: Gerali et al. (2010); US with 3D: Clerc et al. (2015); Norway-NEMO: Kockerols, Kravik and Mimir (2021).
¹Under moderate crisis probability and severity. ²Under higher crisis probability and severity.

To make a quantitative assessment, we can think of the lower effectiveness of hybrid instruments in absorbing losses as representing a lower effective capital ratio. This is the reason why we implement solvency regulation in terms of two quanta of additional capital requirements: 2.5 and 5 percentage point increases which take place gradually over 20 quarters. The 5 percentage point scenario is broadly in line with the actual implementation of Basel III when the quality of capital is included and under the assumption that hybrid instruments were completely ineffective at absorbing losses. The 2.5 percentage point scenario corresponds to the view that quality was not important because the original buffers were effective. We do not take a stance on which of the above extreme assumptions represents the truth, but they do encompass the range of possibilities. Possible non-linearities may lead to responses that are not proportional, hence the need to consider both scenarios explicitly.

The Basel III policy is implemented in the following way in the models. We calibrate the models to data in 2016, i.e., focusing on a time when the transition to the new policy regime is over. We match banks' capital ratios, their probability of failure (for those models where this variable is meaningful) as well as other macro and financial variables. We then perform a counterfactual experiment: we ask the question of how the economy and the financial system would have looked if bank minimum capital ratios had not risen 2.5 percent (or 5 percent) as a result of the Basel III reforms. We measure the impact of Basel III by looking at some simple magnitudes that are usually associated with the benefits and costs of regulation. Table 4 exhibits the results for each jurisdiction. Negative numbers indicate that a variable has declined as a result of Basel III (e.g., bank default probability).

Table 4 suggests that Basel III has been a successful policy. There are some differences across models which can be traced back to differences in the scope of the assessment and in the transmission channels of regulation. But in most models, which include both the costs and benefits of regulation, the long-run effects of Basel III are positive on GDP. This is the case for the 3D model applied to the euro area and the United States, as well as the model by de Bandt and Chahad (2016) applied to the euro area.

The versions of the 3D model employed by the ECB and the Board of Governors exhibit a positive effect of Basel III on GDP

in the long run (1.19 percent in the euro area, 0.87 percent in the United States), associated with a slowdown in the short run. Table 4 shows that the implementation of higher capital requirements leads to a significant reduction in the probability of bank failure (-7.5 percentage points in the euro area and -9.21 percentage points in the United States). Lending spreads increase in all countries. It is worth investigating further the channels that drive these results. The 3D model places banks and their solvency at the heart of the model. All credit (to firms and households) is bank intermediated and the state of bank balance sheets matters greatly for the cost of funds for final borrowers. Since banks are competitive in the model, loan interest rates reflect the weighted average of the cost of equity and the real cost of bank debt (including insured and uninsured debt liabilities) where the weight on equity is the bank's capital ratio. This suggests that changing capital requirements affect the supply of loans through two main channels.

First of all, a key cost of higher capital ratios arises due to the assumption that the risk-adjusted expected return on equity (i.e., the cost of equity) is higher than that on short-term deposits due to limited participation in the equity market. So a bank that funds itself more with equity has a higher weighted average cost of funds. The model also assumes that the required returns on equity do not change as banks' capital ratios increase. In other words, we do not have the so-called Modigliani-Miller offset. These assumptions are in line with the approach in much of the literature (Birn et al. 2020). The assumptions are driven by the difficulty of macroeconomic models to match risk premia in financial markets (the so-called equity premium puzzle). In addition, judging from the fact that market-to-book ratios in banking stocks have significantly declined since 2007, it is difficult to believe that banks nowadays face a lower cost of issuing equity compared to the pre-crisis period. However, the impact of assuming costly equity on banks' overall weighted average cost of funds is quantitatively small: with an 8 percent equity premium, a 1 percentage point change in the capital ratio would increase the weighted average cost of funds by 8 basis points. Nevertheless, assuming a lower cost of equity or a Modigliani-Miller offset would reduce the costs of capital measures further.

Second, and going in the opposite direction, a higher capital ratio reduces the weighted average private cost of funds due to its effect

on the cost of uninsured debt. In the 3D model, market discipline from the uninsured debt market is assumed to be weak due to the opacity of bank balance sheets. Strict capital regulation together with effective supervision is therefore a key mechanism, which reassures uninsured debt investors that banks are safe. As regulation limits the scope for bank risk-taking and moral hazard, interest rate on bank debt falls, exerting an expansionary effect on credit supply. This effect is strong when the probability of failure of banks is relatively high, and it can dominate the contractionary effect of shifting the composition of liabilities towards costlier equity. This explains why the 3D model for the euro area and United States produces an increase in lending and GDP. Given the estimated risks faced by banks, the models imply that the banking system would have been in a very weak state if pre-crisis capital ratios had persisted to the present. This would have affected its ability to supply credit to the real economy with negative effects on real economic activity. Making banks safe is therefore not a cost to the economy but a precondition for a well-functioning credit market.

Finally, some of the costs of financial instability are met by taxpayers who provide implicit or explicit safety net guarantees to bank debt or deposit holders. Higher capital ratios reduce also these public costs: this is reported in Table 4 in the column named “Bank Bailout Costs as a % of GDP.” All of the above private and social benefits of capital regulation are quantitatively very large in the 3D model calibrated to the euro area and the United States.

There are two other models calibrated to the euro area in Table 4. First, the model of de Bandt and Chahad (2016) also shows a positive effect on GDP, although the magnitude is smaller. It is both interesting and instructive to think through the mechanisms in the de Bandt and Chahad (2016) framework and why the benefits appear to be smaller. The model features banks which face random exogenous withdrawals in the spirit of Angeloni and Faia (2013). If the liquidation value of the bank’s assets is smaller than the value of demand deposits, there exists a self-fulfilling “run equilibrium” in the spirit of Diamond and Dybvig (1983). The run can be prevented by either choosing to fund with equity which cannot be withdrawn or by investing in a large amount of liquid assets which do not incur liquidation costs. Hence the probability of crisis depends on the bank’s capital ratio as well as its holdings of liquid assets (this is the aim

of the liquidity coverage ratio). What the results in Table 4 show is that the capital ratio has a positive but relatively small effect on the probability of a run on the bank. Indeed, this is intuitive. If the illiquid assets lead to large losses in a run and loan fire-sale scenario, the bank needs extremely large capital buffers in order to survive. A relatively small capital increase like the one we consider helps to some extent but cannot completely remove the inherent fragility of the fractional reserve banking system. This is why Basel III also introduced liquidity requirements which increase the ability of banks to meet withdrawals without having to liquidate illiquid assets at a loss. In subsequent analysis (see Table 5 in Section 4.2) we will use the de Bandt and Chahad (2016) model to show that both capital and liquidity regulation have an important role to play in making banks safer from runs.

Second, the regime-switching version of Norges Bank's main DSGE model (NEMO), is simulated with two different sets of assumptions on the link between household credit growth and financial crises. In the case of moderate sensitivity of crisis probability and severity to five-year real household credit growth, Basel III has a small negative effect on GDP although it reduces both the crisis probability and the severity. The small negative overall effect on GDP occurs because the reduction in real economic activity in normal times slightly outweighs the gain from the reduction in the probability of financial crisis. In this scenario, costs and benefits of Basel III seem finely balanced. However, when the sensitivity of both probability and severity to credit growth is doubled, Basel III has large positive effects on GDP and its net benefits become substantial. In particular, the negative impact on GDP turns into a positive effect as higher requirements help reduce the probability and the severity of a deeper financial crisis (about 10 percent reduction in output during the crisis). In the latter case, the ergodic mean of GDP increases by 2.1 percent in the long run under the higher capital requirements of the Basel III regime. In terms of mechanisms in this framework, the costs of financial regulation are linked to the excess cost of equity relative to debt. Similarly to the 3D model, a higher capital ratio increases the weighted average cost of capital for banks in normal times, hence reducing economic activity since banks need to accumulate larger net worth by charging higher lending rates. The benefits of higher capital requirements arise because

of a reduced frequency and severity of financial crises. When a crisis occurs, the economy switches from a normal times regime to a crisis regime where macroeconomic and financial shocks are asymmetrically large (having higher volatility) and some structural model parameters regarding housing and banking sectors switch.⁶ This leads to amplified adverse financial accelerator effects leading to private sector deleveraging with substantial negative effects on economic activity.


All in all, the results of the models crucially depend on the assumptions regarding the magnitude and the sensitivity of the bank default probability (euro area and United States) or financial crisis probability (Norway) to underlying features of the banking sector. This is consistent with the LEI study (BCBS 2010) and Birn et al. (2020).

In additional exercises, we also assess the opportunity costs related to the implementation of Basel III without separate consideration of benefits. The Gerali et al. (2010) framework for the euro area (“cost approach”), which only identifies the cost of implementation of the regulation, yields a negative effect on GDP, but this result is an obvious consequence of not modeling the benefits of regulation. The long-run net benefits of the Basel III framework could be estimated by comparing the steady-state increase in GDP in the euro-area 3D model (1.2 percent) or in the model by de Bandt and Chahad (2016) (0.2 percent), with the decrease in GDP for the euro area according to the Gerali et al. (2010) model (GDP growth down by 0.4 percent). This yields a long-run net benefit between 0.6 and 1.6 percent of GDP in the euro area.

Finally we investigate the transition dynamics between Basel II and III. The models employed by the European Central Bank and the Board of Governors exhibit a positive long-run effect of Basel III on GDP although there is a temporary slowdown over the transition, as shown in Figure 1 for the 3D euro area model (transitions for the four other models are displayed in the appendix). When capital requirements are increased, the available bank capital can support less lending. Financial institutions therefore contract credit supply

⁶A crisis can occur any point in time in the model given the estimated crisis probability that depends on five-year cumulative real household credit growth in the model.

Figure 1. Transition from 14 Percent Capital Ratio to 16.5 Percent in the Euro Area with 3D Model



Note: Variables are expressed in deviation from initial steady state. “3D Model” refers to the model used by Mendicino et al. (2020).

and lending margins increase. This raises bank profits, and bank capital (“total capital” in the figure) starts to rise. In the long run, total lending recovers as bank capital is accumulated by retained earnings helping to support more loans.

Over the first 10 quarters or so, however, the reduction in lending depresses the investment of bank-dependent firms. This reduces aggregate GDP and leads to a small decline in inflation and a loosening of monetary policy. The monetary accommodation boosts consumption and the investment of firms that do not depend on bank funding (not shown in the figure).⁷ However, the increase in these components of demand is not enough to offset the fall in investment

⁷The 3D model includes a bank-independent sector which funds itself directly from the household sector. Here we have in mind larger firms that can issue equity or corporate bonds.

demand from bank-dependent firms and output declines. In the long run, once capital has fully adjusted, GDP increases.

There are a number of factors which determine the size of the transitional costs from an increase in bank capital requirements. As discussed in Mendicino et al. (2020), the monetary policy reaction is important. When the policy rate is reduced aggressively at the start of the implementation period, this helps to counteract the negative effects on aggregate demand from the squeeze on the credit availability for smaller bank-dependent firms. At the effective lower bound (ELB) on nominal interest rate, this may require the use of non-standard measures to ensure that real economic activity does not suffer too much. When it is difficult to deliver a sizable monetary stimulus, a more gradual implementation period may be optimal. From this perspective, the long time period given to banks to comply with Basel III seems like a good policy which would have significantly smoothed the transitional costs of the increase in bank capital requirements.

Darracq-Pariès et al. (2022) also emphasize several other factors which may affect the size of the transitional costs. First, the presence of large voluntary buffers can help banks smooth the increase in bank capital requirements even when they are imposed without a long implementation period. Second, the ability and willingness of banks to issue equity or cut dividends also reduces the transitional costs. Finally, the ability of non-bank funding sources to expand over the transition is also important.⁸ In the aftermath of the GFC, voluntary buffers were depleted for many banks and the costs of issuing equity were high. Hence, the long implementation period for Basel III was therefore important in ensuring that the transitional costs were manageable for the real economy.

4.2 *Liquidity Scenario*

There are two liquidity instruments in Basel III—the liquidity coverage ratio (LCR) and the net stable funding ratio (NSFR). The

⁸Non-bank funding sources are of course a double-edged sword. They can reduce the costs of the transition to higher capital requirements but they may also reduce the effectiveness of prudential measures due to “leakages.”

LCR involves the obligation to hold a sufficient quantity of high-quality liquid assets (HQLA) so as to withstand one month of elevated deposit withdrawals. The NSFR involves the obligation to fund long-term assets at least in part with longer-term liabilities (e.g., bank bonds). In practice, both the LCR and NSFR are complex regulations, which aim to increase banks' resilience to funding stress. To make the analysis operational, we concentrate on the LCR and provide the results of two sets of models in Table 5: (i) analysis of the costs of liquidity following the approach in Hoerova et al. (2018); and (ii) a more detailed analysis based on the multi-asset model of de Bandt and Chahad (2016) with bank runs.

Regarding the first approach, most models do not include an analysis of the benefits of the LCR. In this case, the simulations follow Hoerova et al. (2018) and only consider the impact of an LCR scenario through its effect on bank profits, measuring the opportunity cost of raising additional deposits and investing in lower-yielding HQLA. Regulation affects the profit and loss statement, as the return on HQLA is lower than interest and non-interest costs on deposits needed to fund the HQLA holdings. Hoerova et al. (2018) identify the cost of holding a unit of HQLA to be 0.68 percent, meaning that a bank makes a loss of EUR 0.68 on an HQLA holding worth EUR 100, which is fully financed with deposits.⁹ They argue that the move from pre-crisis LCR levels to full compliance with the new Basel III standard (100 percent LCR) involves banks increasing their HQLA holdings by an amount worth 10 percent of total deposits.¹⁰ In equilibrium, loan rates must increase following a negative shock to loan supply in order to restore banks' profitability at least partially. This is how the LCR exerts a negative impact on lending and economic activity.

A crucial caveat is in order. We assume that HQLA are government bonds with a zero capital risk weight. If the HQLA have a non-zero risk weight or if the leverage ratio is the binding constraint,

⁹For the Norwegian banks, the cost of holding a unit of HQLA is calculated to be 0.46 percent.

¹⁰For the Norwegian banks, implementing an LCR of 100 percent is approximated by asking the bank to hold government bonds and covered bonds equal to 11.2 percent of deposits.

the LCR will also directly crowd out lending and will lead to larger costs than those considered here.

The only framework that quantifies the benefits of liquidity regulation is the model by de Bandt and Chahad (2016), where the LCR requirement is explicitly modeled in a multi-asset framework, mimicking the actual regulation, which allows an assessment of the impact of liquidity regulation on the probability of runs.

Table 5 shows the steady-state impact of the LCR regulations.¹¹ For the euro area 3D model, the implementation of the LCR does not affect banks' probability of default (PD), and consequently bailout costs and private lending spreads are also unaffected.¹² As already discussed, the LCR regulation affects bank profitability negatively and, in a partial equilibrium setting, raises the default probabilities of banks. In general equilibrium, however, following a negative shock on loan supply, banks increase their lending rates. Thus, bank solvency does not suffer, but as the fourth column of Table 5 shows, lending spreads over bank funding costs increase by 6 basis points in the case of the LCR for the euro area with the 3D model, and 20 basis points with the model by de Bandt and Chahad (2016). We also notice that the higher cost of funding for borrowing firms and households reduces real economic activity by a moderate amount. The LCR reduces consumption by 0.1 percent, investment by 0.31 percent, and GDP by 0.14 percent in the new steady state. Total lending falls by 0.73 percent. These relatively small costs should be set against the benefits of the regulatory measures.

The model by de Bandt and Chahad (2016) is the only one that finds a positive long-run effect (GDP rises by 0.27 percent), due to a strong expectation channel associated with a lower probability of bank runs which makes banks safer (bank PDs fall by 0.68 percent). When depositors expect a lower probability of a bank run, their expected returns from holding bank deposits increase and they demand a lower interest rate. This increases banks' profitability,

¹¹In the case of the de Bandt and Chahad model, the simulation includes both solvency and liquidity requirements.


¹²There is a very small increase in banks' PDs, but this is less than 1 basis point and is rounded to zero in the table. In the transitional figures, a negative but extremely small effect on bank solvency can be seen.

allowing them to accumulate more capital through retained earnings and supporting more lending, which rises by almost 2 percent. It is worth delving with a little more detail into the results from this particular model in Table 5 and comparing them to those in Table 4. Note that for that model the scenario in Table 5 includes an increase in both the solvency and liquidity requirements in order to mimic the actual implementation of Basel III. We see that adding the LCR ratio (in Table 5) generates a further increase in GDP, a bigger increase in consumption (which is arguably a better measure of household welfare), and a bigger decline in banks' probability of failure. This finding confirms the intuition that each risk should be targeted with an appropriate policy tool. The de Bandt and Chahad (2016) model emphasizes that a liquidity policy tool (in this case the LCR) is the most effective way to mitigate liquidity risk. Capital tools also help with liquidity risk, but they are, relatively speaking, less effective. In contrast, as the 3D model results for the euro area and the United States in Table 4 showed, the capital ratio is extremely effective at mitigating solvency risk.

We end our analysis of the impact of the LCR by including the transitional dynamics from a world with no liquidity regulation to a world with full compliance with the LCR. The regulation requires banks to increase the quantity of liquid assets on their balance sheets. In addition, the regulation stimulates banks to increase the quality of their HQLA: this is because the LCR gives a higher weight to government bonds compared to corporate debt. The change in the composition of HQLA carries a small cost because, although government bonds have a lower return than corporate bonds, they have a bigger negative effect on the probability of a bank run.

Figure 2 shows how banks react to the policy: initially they switch the composition, and government debt holdings jump sharply to satisfy the LCR while corporate bond holdings decline. The higher quantity and quality of liquid securities reduce the probability of a bank run. This, in turn, lowers risk premia on bank debt, raises bank profitability, and leads to the accumulation of more capital through retained earnings. As a result, lending, investment, consumption, and GDP gradually increase. As lending rises, banks continue to grow their HQLA (both government and corporate debt) over time in order to continue satisfying the higher LCR. Importantly, lending increases to both small and large firms. In addition, the reduction

Figure 2. Impact of LCR Implementation with the Model by de Bandt and Chahad (2016) for the Euro Area (with regulated and effective capital ratios)



Note: This figure represents the dynamics of several macroeconomic and financial variables to an increase in LCR and capital ratio. The top left two charts exhibit the actual (or effective) ratios contrasted to the regulatory requirements. All variables are expressed as deviations from the initial steady state.

in bank run probability allows an increase in deposits, in contrast to what is observed when only solvency requirements are increased.

Overall, the results in this section suggest that the costs of the imposition of the LCR in the Basel III reforms were very modest while the benefits were significant especially in reducing the liquidity risk faced by the banking system.

5. Conclusion

In this paper, we survey channels of transmission of prudential regulation proposed in the literature since the global financial crisis. We also conduct harmonized simulation exercises examining the effects of Basel III regulation, based on off-the-shelf DSGE models used in various jurisdictions. Our conclusion is that Basel III leads to an increase in GDP, although some models show some negative effects. The increase in GDP is achieved through a reduction in the cost of borrowing, which, in turn, is achieved through a reduction in cost of uninsured debt as higher capital requirement reduces the probability of bank failures, leading to a net increase in credit supply.


However, the models only offer a partial assessment of the macro-economic impact of the new regulatory environment. In particular, when assessing the effects of banking regulations, it is crucial to distinguish models that assess both costs and benefits (e.g., the 3D model and NEMO) from models that only assess costs (Gerali et al. 2010). A few limitations have been identified: the models are still quite stylized with only one capital variable, namely aggregate capital; funding liquidity has only been incorporated in a basic way; and liquidity regulation is not fully integrated in most models.

Our study also identifies directions in which policy models should expand. Quantitative DSGE models basically focus on capital requirements, while empirical models lack micro-foundations, which is problematic for policy analysis. More complex issues, such as interactions between multiple regulations, still depend on qualitative models. In addition, these models have only recently started to investigate additional policy issues, and there is scope for further research regarding the role of shadow banking and the interaction between unconventional monetary policy and financial stability policy.

Further research would also include investigating further the micro-foundations of different types of financial crises (bank based, market based, asset price bubbles, etc.), the role of rational versus non-rational expectations, and the distributional effects of financial regulation.


Appendix

Figure A.1. Transition from 14 Percent Capital Ratio to 16.5 Percent in the Euro Area (de Bandt and Chahad)




Note: Variables are expressed as deviations from the initial steady state.

Figure A.2. Transition from 14 Percent Capital Ratio to 16.5 Percent in the Euro Area (NEMO)




Note: Variables are expressed as deviations from the initial steady state.

Figure A.3. Transition from 14 Percent Capital Ratio to 16.5 Percent in the Euro Area (Gerali et al.)



Note: Variables are expressed as deviations from the initial steady state.

Figure A.4. Transition from 14 Percent Capital Ratio to 16.5 Percent in the United States (3D)



Note: Variables are expressed as deviations from the initial steady state.

References

- Adrian, T., and N. Boyarchenko. 2018. "Liquidity Policies and Systemic Risk." *Journal of Financial Intermediation* 35 (Part B): 45–60.
- Adrian, T., F. Duarte, N. Liang, and P. Zabczyk. 2020. "NKV: A New Keynesian Model with Vulnerability." *AEA Papers and Proceedings* 110 (May): 470–76.
- Angeloni, I., and E. Faia. 2013. "Capital Regulation and Monetary Policy with Fragile Banks." *Journal of Monetary Economics* 60 (3): 311–24.
- Badarau, C., and G. Levieuge. 2011. "Assessing the Effects of Financial Heterogeneity in a Monetary Union a DSGE Approach." *Economic Modelling* 28 (6): 2451–61.

- Basel Committee on Banking Supervision. 2010. "An Assessment of the Long-Term Economic Impact of Stronger Capital and Liquidity Requirements." Document No. 173, Bank for International Settlements.
- . 2012. "Models and Tools for Macroprudential Analysis." Working Paper No. 21, Bank for International Settlements.
- . 2021. "Assessing the Impact of Basel III: Evidence from Macroeconomic Models: Literature Review and Simulations." Working Paper No. 38, Bank for International Settlements.
- Begenau, J. 2020. "Capital Requirements, Risk Choice, and Liquidity Provision in a Business-Cycle Model." *Journal of Financial Economics* 136 (2): 355–78.
- Begenau, J., and T. Landvoigt. 2022. "Financial Regulation in a Quantitative Model of the Modern Banking System." *Review of Economic Studies* 89 (4): 1748–84.
- Bennani T., C. Couaillier, A. Devulder, S. Gabrieli, J. Idier, P. Lopez, T. Piquard, and V. Scalone. 2017. "An Analytical Framework to Calibrate Macroprudential Policy." Working Paper No. 648, Banque de France.
- Bernanke, B. 2005. "The Global Saving Glut and the US Current Account Deficit." Speech, Board of Governors of the Federal Reserve System.
- Bernanke, B., 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). North-Holland.
- Birn, M., O. de Bandt, S. Firestone, M. Gutiérrez Girault, D. Hancock, T. Krogh, H. Mio, D. P. Morgan, A. Palvia, V. Scalone, M. Straughan, and A. Uluc. 2020. "The Costs and Benefits of Bank Capital — A Review of the Literature." *Journal of Risk and Financial Management* 13 (4): 1–25.
- Boissay, F., and F. Collard. 2016. "Macroeconomics of Bank Capital and Liquidity Regulations." Working Paper No. 596, Bank for International Settlements.
- Boissay, F., F. Collard, and F. Smets. 2016. "Booms and Banking Crises." *Journal of Political Economy* 124 (2): 489–538.
- Brunnermeier, M., and Y. Sannikov. 2014. "A Macroeconomic Model with a Financial Sector." *American Economic Review* 104 (2): 379–421.

- 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.
- 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, June): 9–63.
- Coimbra, N., and H. Rey. 2023. "Financial Cycles with Heterogeneous Intermediaries." Forthcoming in *Review of Economic Studies*.
- Collard, F., H. Dellas, B. Diba, and O. Loisel. 2017. "Optimal Monetary and Prudential Policies." *American Economic Journal: Macroeconomics* 9 (1): 40–87.
- Corbae, D., and P. D'Erasmus. 2019. "Capital Requirements in a Quantitative Model of Banking Industry." NBER Working Paper No. 25424.
- 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.
- Crouzet, N. 2018. "Aggregate Implications of Corporate Debt Choices." *Review of Economic Studies* 85 (3): 1635–82.
- Darracq-Pariès, M., P. Karadi, C. Kok, and K. Nikolov. 2022. "The Impact of Capital Requirements on the Macroeconomy: Lessons from Four Macroeconomic Models of the Euro Area." *International Journal of Central Banking* 18 (5, December): 197–246.
- Davis, S. 2010. "The Adverse Feedback Loop and the Effects of Risk in Both the Real and Financial Sectors." Globalization Institute Working Paper No. 66, Federal Reserve Bank of Dallas.
- de Bandt, O., and M. Chahad. 2016. "A DSGE Model to Assess the Post Crisis Regulation of Universal Banks." Working Paper No. 602, Banque de France.
- de Bandt, O., S. Lecarpentier, and C. Pouvelle. 2021. "Determinants of Banks' Liquidity: A French Perspective on Interactions between Market and Regulatory Requirements." *Journal of Banking and Finance* 124 (March): Article 106032. <https://doi.org/10.1016/j.jbankfin.2020.106032>.

- De Fiore, F., and H. Uhlig. 2011. "Bank Finance versus Bond Finance." *Journal of Money, Credit and Banking* 43 (7): 1399–1421.
- . 2015. "Corporate Debt Structure and the Financial Crisis." *Journal of Money, Credit and Banking* 47 (8): 1571–98.
- De Nicolò, G., A. Gamba, and M. Lucchetta. 2014. "Microprudential Regulation in a Dynamic Model of Banking." *Review of Financial Studies* 27 (7): 2097–2138.
- Diamond, D. W., and P. H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91 (3): 401–19.
- Dou, W. W., X. Fang, A. W. Lo, and H. Uhlig. 2021. "Macro-Finance Models with Nonlinear Dynamics." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3457590.
- Durdu, C. B., and M. Zhong. 2023. "Understanding Bank and Nonbank Credit Cycles: A Structural Exploration." *Journal of Money, Credit and Banking* 55 (1): 103–42.
- Elenev, V., T. Landvoigt, and S. Van Nieuwerburgh. 2021. "A Macroeconomic Model with Financially Constrained Producers and Intermediaries." *Econometrica* 89 (3): 1361–1418.
- Fève P., A. Moura, and O. Pierrard. 2019. "Shadow Banking and Financial Regulation: A Small-Scale DSGE Perspective." *Journal of Economic Dynamics and Control* 101 (April): 130–44.
- 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.
- Gertler, M., and N. Kiyotaki. 2015. "Banking, Liquidity and Bank Runs in an Infinite Horizon Economy." *American Economic Review* 105 (7): 2011–43.
- Gertler, M., N. Kiyotaki, and A. Prestipino. 2016. "Wholesale Banking and Bank Runs in Macroeconomic Modeling of Financial Crises." In *Handbook of Macroeconomics*, Vol. 2B, ed. J. B. Taylor and H. Uhlig, 1345–1425 (chapter 16). North-Holland.
- . 2020. "A Macroeconomic Model with Financial Panics." *Review of Economic Studies* 87 (1): 240–88.

- Goldstein, I., and A. Pauzner. 2005. "Demand-Deposit Contracts and the Probability of Bank Runs." *Journal of Finance* 60 (3): 1293–1327.
- Hart, O., and J. Moore. 1994. "A Theory of Debt Based on the Inalienability of Human Capital." *Quarterly Journal of Economics* 109 (4): 841–79.
- He, Z., and A. Krishnamurthy. 2019. "A Macroeconomic Framework for Quantifying Systemic Risk." *American Economic Journal: Macroeconomics* 11 (4): 1–37.
- Hirakata, N., N. Sudo, and K. Ueda. 2011. "Do Banking Shocks Matter for the U.S. Economy?" *Journal of Economic Dynamics and Control* 35 (12): 2042–63.
- Hoerova, M., C. Mendicino, K. Nikolov, G. Schepens, and S. Van den Heuvel. 2018. "Benefits and Costs of Liquidity Regulation." Working Paper No. 2169, European Central Bank.
- Holden, T. D., P. Levine, and J. M. Swarbrick. 2020. "Credit Crunches from Occasionally Binding Bank Borrowing Constraints." *Journal of Money, Credit and Banking* 52 (2–3): 549–82.
- Holmstrom, B., and J. Tirole. 1997. "Financial Intermediation, Loanable Funds, and the Real Sector." *Quarterly Journal of Economics* 112 (3): 663–91.
- Ikeda, D. 2018. "Bank Runs, Prudential Tools and Social Welfare in a Global Game General Equilibrium Model." Working Paper No. 732, Bank of England.
- Irani, R. M., R. Iyer, R. R. Meisenzahl, and J.-L. Peydró. 2021. "The Rise of Shadow Banking: Evidence from Capital Regulation." *Review of Financial Studies* 34 (5): 2181–2235. doi.org/10.1093/rfs/hhaa106.
- Jeanne, O., and A. Korinek. 2020. "Macroprudential Regulation versus Mopping up after the Crash." *Review of Economic Studies* 87 (3): 1470–97.
- Jiménez, G., S. Ongena, J.-L. Peydró, and J. Saurina. 2017. "Macroprudential Policy, Countercyclical Bank Capital Buffers, and Credit Supply: Evidence from the Spanish Dynamic Provisioning Experiments." *Journal of Political Economy* 125 (6): 2126–77.
- Jondeau, E., and J. G. Sahuc. 2022. "Bank Capital Shortfall in the Euro Area." *Journal of Financial Stability* 62 (October): Article 101070.

- Kale, J., and C. Meneghetti. 2011. "The Choice Between Public and Private Debt: A Survey." *IIMB Management Review* 23 (1): 1–66.
- Kashyap, A. K., D. P. Tsomocos, and A. P. Vardoulakis. 2020. "Optimal Bank Regulation in the Presence of Credit and Run-Risk." NBER Working Paper No. 26689.
- Kiyotaki, N., and J. Moore. 1997. "Credit Cycles." *Journal of Political Economy* 105 (2): 211–48.
- Kockerols, T., E. M. Kravik, and Y. Mimir. 2021. "Leaning against the Wind with Persistent Financial Cycles and Occasional Crises." Working Paper No. 11, Norges Bank.
- Korinek, A., and A. Simsek. 2016. "Liquidity Trap and Excessive Leverage." *American Economic Review* 106 (3): 699–738.
- Kravik, E. M., and Y. Mimir. 2019. "Navigating with NEMO." Staff Memo No. 5, Norges Bank.
- Lang, J., and M. Forletta. 2020. "Cyclical Systemic Risk and Downside Risks to Bank Profitability." Working Paper No. 2405, European Central Bank.
- Macroeconomic Assessment Group. 2010. "Assessing the Macroeconomic Impact of the Transition to Stronger Capital and Liquidity Requirements." Report, Bank for International Settlements.
- Martinez-Miera, D., and R. Repullo. 2017. "Search for Yield." *Econometrica* 85 (2): 351–78.
- . 2019. "Markets, Banks and Shadow Banks." Working Paper No. 2234, European Central Bank.
- Martinez-Miera, D., and J. Suarez. 2014. "Banks' Endogenous Systemic Risk Taking." CEMFI.
- Meeks, R., B. Nelson, and P. Alessandri. 2017. "Shadow Banks and Macroeconomic Instability." *Journal of Money, Credit and Banking* 49 (7): 1483–1516.
- Meh, C. A., and K. Moran. 2010. "The Role of Bank Capital in the Propagation of Shocks." *Journal of Economic Dynamics and Control* 34 (3): 555–76.
- Mendicino, C., K. Nikolov, J. Suarez, and D. Supera. 2018. "Optimal Dynamic Capital Requirements." *Journal of Money, Credit and Banking* 50 (6): 1271–97.
- . 2020. "Bank Capital in the Short and in the Long Run." *Journal of Monetary Economics* 115 (November): 64–79.

- Miller, S., and R. Sowerbutts. 2018. "Bank Liquidity and the Cost of Debt." Working Paper No. 707, Bank of England.
- Morris, S., and H. S. Shin. 2001. "Global Games: Theory and Applications." Discussion Paper No. 1275R, Cowles Foundation.
- Plantin, G. 2015. "Shadow Banking and Bank Capital Regulation." *Review of Financial Studies* 28 (1): 146–75.
- Schroth, J. 2021. "Macroprudential Policy with Capital Buffers." *Journal of Monetary Economics* 118 (March): 296–311.
- 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.
- Suarez J. 2022. "Growth-at-Risk and Macroprudential Policy Design." *Journal of Financial Stability* 60 (June): Article 101008.
- Swarbrick, J. 2019. "Lending Standards, Productivity and Credit Crunches." Working Paper No. 2019–2025, Bank of Canada.
- Townsend, R. M. 1979. "Optimal Contracts and Competitive Markets with Costly State Verification." *Journal of Economic Theory* 21 (2): 265–93.
- Urban, J. 2019. "Credit Cycles Revisited." <https://ssrn.com/abstract=3310127> or <http://dx.doi.org/10.2139/ssrn.3310127>.
- Van den Heuvel, S. 2006. "The Bank Capital Channel of Monetary Policy." Meeting Papers No. 512, Society for Economic Dynamics.
- . 2019. "The Welfare Effects of Bank Liquidity and Capital Requirements." *Journal of Monetary Economics* 55 (2): 298–320.
- Walther, A. 2016. "Jointly Optimal Regulation of Bank Capital and Liquidity." *Journal of Money, Credit and Banking* 48 (2–3): 415–48.

The Impact of SNB Monetary Policy on the Swiss Franc and Longer-Term Interest Rates*

Fabian Fink, Lukas Frei, Thomas Maag, and Tanja Zehnder
Swiss National Bank

We estimate the impact of monetary policy announcements by the Swiss National Bank on the Swiss franc and on the expected path of future short-term interest rates. Monetary policy announcement effects are identified using the identification-through-heteroskedasticity approach. The approach accounts for the simultaneous relation of exchange rates and interest rates. We find that from 2000–11, an announcement of a monetary policy tightening appreciated the nominal Swiss franc on the same day. Importantly, the results indicate that simple methods that do not adequately account for simultaneity between exchange rates and interest rates yield biased and typically non-significant estimates. Our findings further suggest that monetary policy announcements affect medium- to longer-term expectations, which in turn influence the Swiss franc.

JEL Codes: E52, E58, E43, F31.

*We would like to thank Angela Abbate, Benjamin Anderegg, Daniele Ballinari, Lucas Fuhrer, Christian Grisse, Oliver Gloede, Basil Guggenheim, Christian Hepenstrick, Matthias Juettner, Carlos Lenz, Peter Kugler, Andrea Maechler, Thomas Moser, Silvio Schumacher, Barbara Stahel, Marcel Zimmermann, and the seminar participants at the August 2019 SNB brown bag workshop for valuable comments and suggestions. We also thank Fernanda Nechio, the editor, as well as the anonymous referee. Moreover, we would like to thank the entire FX Trading Strategy and Technology unit for helpful discussions and their provision of the shared data processing and analytics platform. Thanks to Benjamin Brunner who shared his data set on the decomposition of Swiss government bond yields with us. The views, opinions, findings, and conclusions or recommendations expressed in this paper are strictly those of the authors, and they do not necessarily reflect the views of the Swiss National Bank. The Swiss National Bank takes no responsibility for any errors or omissions in, or for the correctness of, the information contained in this paper. E-mail addresses: fabian.fink@snb.ch (corresponding author), lukas.frei@snb.ch, thomas.maag@snb.ch, tanja.zehnder@snb.ch. Swiss National Bank, Boersenstrasse 15, P.O. Box, 8022 Zurich, Switzerland.

1. Introduction

Economic theory suggests that an increase in the policy rate appreciates the currency in the short and long run. This result is explained by two economic principles, namely, interest rate parity and purchasing power parity. Identifying the effects of policy rate changes on exchange rates is challenging due to endogeneity and simultaneity. Exchange rate and interest rate movements reflect not just monetary policy but also other drivers such as safe-haven flows into the Swiss franc.¹ For example, a negative risk shock tends to result in *lower* interest rates and a *stronger* Swiss franc because of safe-haven flows into the Swiss franc and an expected loosening of monetary policy in response to a Swiss franc appreciation. This holds particularly true for a small open economy such as Switzerland, for which the exchange rate is an important determinant of monetary policy transmission and hence for the monetary policy decisionmaking process.² In contrast, a negative domestic economic shock tends to induce *lower* interest rates and a *weaker* Swiss franc, both due to an expected loosening of monetary policy. Hence, the naively observed relation between interest rates and exchange rates will depend on the type of shock that dominates in a given sample period.

This paper investigates the short-term impact of monetary policy announcements made by the Swiss National Bank (SNB) on the Swiss franc exchange rates and on the Swiss yield curve. A monetary policy announcement includes a monetary policy decision about its policy rate and a conditional inflation forecast.³ Furthermore, a forecast for the Swiss economy is provided. Thus, the monetary policy shock we identify reflects these two components: the

¹The safe-haven characteristics of the Swiss franc have been established by numerous contributions. See Fink, Frei, and Gloede (2022) for an overview.

²Nitschka and Mirkov (2016) estimate Taylor rules for Switzerland, augmented by an effective nominal Swiss franc exchange rate. They find that the estimated effect of Swiss franc appreciation on the three-month LIBOR (London interbank offered rate) target rate is highly significant and negative, reflecting the stabilizing effect of the Swiss National Bank policy. Importantly, Nitschka and Mirkov (2016) report that professional survey expectations anticipate a response of the central bank.

³The forecast is based on the assumption that the policy rate set when the forecast is published will remain constant over the entire three-year forecast horizon.

change of monetary policy and the change of the outlook of the Swiss economy. We identify the causal effect on a daily level using the identification-through-heteroskedasticity (IH) methodology following Rigobon (2003) and Rigobon and Sack (2004). The approach allows for simultaneous (intraday) feedback between interest rates and exchange rates and imposes weaker identifying assumptions than event-study-based approaches.

The IH methodology is based on the fact that the variance of the interest rate shocks is higher on days of monetary policy announcements than on other days. To estimate the effects of monetary policy announcements, we consider the historical time frame spanning from January 1, 2000 to August 31, 2011. During this time, the three-month CHF LIBOR was the SNB's main policy instrument. Our sample includes two policy rate cycles, covering 56 SNB monetary policy announcements and 26 changes in the three-month CHF LIBOR target range.

The main findings can be summarized as follows. First, a monetary policy announcement shock which increases the policy rate leads to a nominal appreciation of the Swiss franc on the same day. The null hypothesis that monetary policy announcement shocks that change the policy rate do not affect the exchange value of the Swiss franc is clearly rejected. Our finding that monetary policy announcement shocks by the SNB are highly relevant for the exchange rate of the Swiss franc is robust to the use of alternative specifications. Second, we provide empirical evidence that simple econometric methods yield biased and typically non-significant results. Not adequately identifying the causal effects may lead to the incorrect conclusion that monetary policy announcements do not affect the Swiss franc. Third, monetary policy announcements affect the curve of expected average future short-term interest rates, which in turn influence the valuation of the Swiss franc.

Our paper contributes to the literature as follows. First, we extend the rather scant empirical literature on Switzerland. In particular, we discuss how to adequately identify causal effects, which is particularly relevant for the Swiss case. Moreover, we contribute to the understanding of transmission by estimating the effects on the expected average future short-term interest rates derived using the Adrian, Crump, and Moench (2013) methodology. Second, with regards to the IH literature that builds on the seminal

contributions of Rigobon (2003) and Rigobon and Sack (2004), we advance the estimation methodology by *jointly* estimating the effects on the EUR/CHF and USD/CHF exchange rates using the generalized method of moments (GMM).⁴ Joint estimation is more efficient and facilitates additional robustness tests.

The analysis is structured as follows. Section 2 summarizes the relevant literature, focusing on evidence for Switzerland. Section 3 outlines the simultaneity issue and our identification approach. Section 4 describes the data. Section 5.1 presents the results on the exchange rate response to a monetary policy announcement shock. In Section 5.2, we discuss the effects on the expected average future short-term interest rates. Section 6 concludes.

2. Related Literature

A large body of the empirical literature investigates the effects of monetary policy surprises on asset prices. As highlighted by Rigobon and Sack (2004), this relation is not only important for financial market participants but also for central banks because asset prices play an important role in monetary policy transmission.

Our contribution relates to the strand of this literature that seeks to identify the effects on exchange rates seen in daily or intra-day data. Causal effects are mostly identified using event studies and by means of instrumental variables or, as in our paper, IH. This strand is distinct from an alternative strand that estimates the effects evident in lower-frequency data, typically using structural vector autoregressions (SVARs).

With regards to the impact of unexpected monetary policy shocks on exchange rates, the international literature reports highly significant and immediate effects (e.g., Kearns and Mannes 2006; Rosa 2011; Ferrari, Kearns, and Schrimpf 2017; Kerksenfischer 2019). The point estimates roughly range from 0.3 percent to 1.5 percent nominal appreciation in response to a contractionary 25 basis point

⁴Exchange rates are quoted in relation to another currency. In this paper we follow market conventions to quote exchange rates, i.e., the first currency is the base currency and the second currency is the price currency. Thus, if you purchase USD/CHF, you would receive one unit of the U.S. dollar in exchange for a payment of Swiss franc.

(bp) interest rate shock. Specifically, considering a panel of seven major central banks, Ferrari, Kearns, and Schrimpf (2017) find that an unexpected policy rate hike of 25 bp causes an immediate appreciation from 1 to 1.5 percent for most of the central banks considered. Focusing on the European Central Bank (ECB), Kerssenfischer (2019) reports that the EUR/CHF exchange rate rises by 0.95 percent in response to a contractionary 25 bp ECB monetary policy shock. Rosa (2011) shows that the surprise components of both the Federal Reserve's (Fed's) monetary policy actions and statements have economically important and highly significant effects on the exchange rate of the U.S. dollar. An unanticipated 25 bp cut in the federal funds target rate is associated on average with a 0.5 percent depreciation of the exchange value of the U.S. dollar, also towards the Swiss franc. Kearns and Manners (2006) investigate the impact of monetary policy on the exchange rate using an event study with intraday data for four countries. An unanticipated tightening of 25 bp leads to a rapid appreciation of approximately 0.35 percent from 1993 to 2004.

In line with international evidence, a rather small literature on effects of SNB policy rate changes confirms that the Swiss franc significantly appreciates in response to contractionary monetary policy shocks. The magnitude of the point estimates is rather broad, however, ranging from less than 0.2 percent to more than 6 percent appreciation in response to a monetary policy announcement of a 25 bp interest rate increase.

The lower bound of this range is set by Ranaldo and Rossi (2010). Using an event-study approach to identify the intraday effects of SNB monetary policy decisions, the authors find that an unexpected 25 bp increase in the three-month CHF LIBOR caused the Swiss franc to appreciate by 0.17 percent towards the U.S. dollar from 2000 to 2005. The upper bound is given by the estimates of Ferrari, Kearns, and Schrimpf (2017), who report that an SNB policy announcement that generates a 25 bp increase in the one-month CHF overnight index swap (OIS) rate caused the Swiss franc to appreciate by 6.25 percent towards the U.S. dollar from September 2010 to September 2015. The magnitude of this effect exceeds what they find for most other central banks, and they note that the findings for Switzerland should be interpreted with care due to the short sample period.

The recent results of Grisse (2020) and Kugler (2020) fall within this range. Similar to our contribution, both papers consider the period 2000–11. Using an instrumental-variable approach and daily data, Kugler (2020) reports that the Swiss franc appreciates by 0.93 percent in response to an unexpected 25 bp increase in the three-month CHF LIBOR. Grisse (2020) estimates the effects using a weekly SVAR and identifies monetary policy shocks based on the co-movement of interest rates and stock prices: A contractionary 25 bp shock causes the Swiss franc to appreciate by 1.0 percent against the euro and by 0.75 percent against the U.S. dollar in the same week.⁵

3. Empirical Model

This section describes first the simultaneity problem when analyzing the effect of monetary policy changes on exchange rates. This section then presents the model, the identification strategy, and the estimation methodology.

3.1 *Model*

When modeling the response of exchange rates to interest rates, we need to take into account the simultaneous effect that a change in the exchange rate may have on the interest rate. Today's exchange rate depends, among other factors, on the expected path of the interest rate, which determines the relative attractiveness of investments in Swiss franc. However, the interest rate is affected by the expected monetary policy response to exchange rate variations. For

⁵Related research confirms the relevance of policy rates for the Swiss franc. Lenz and Savioz (2009) analyze the determinants of the Swiss franc exchange rate against the euro. They find that Swiss monetary policy contributed between 7 percent and 15 percent to variations of the exchange rate from 1981 to 2007. Rudolf and Zurlinden (2014) find an impact of approximately 0.2 percent in an estimated DSGE model for the period 1983–2013. Based on a calibrated dynamic stochastic general equilibrium (DSGE) model of the Swiss economy, the results of Cuche-Curti, Dellas, and Natal (2009) point towards 0.25 percent appreciation for a restrictive 25 bp policy rate shock in the period 1975 to 2006. Canetg and Kaufmann (2019) analyze the impact of SNB's debt security auctions from 2008 to 2011 on financial market variables. They identify a money market as well as an expectation shock and find that the two shocks explain up to 80 percent of the forecast-error variance of the Swiss franc.

a small open economy such as Switzerland, market participants anticipate that the exchange rate is an important factor in the monetary policy decisionmaking process. The latter is in line with the augmented Taylor-rule estimates in Nitschka and Mirkov (2016).

We assume that the change in the policy rate Δi_t and the change in the nominal exchange rate Δs_t are described by the following simultaneous equation system:

$$\text{Exchange rate response function} \quad \Delta s_t = \alpha \Delta i_t + \gamma_s z_t + \eta_t, \quad (1)$$

$$\text{Interest rate response function} \quad \Delta i_t = \beta \Delta s_t + \gamma_i z_t + \varepsilon_t, \quad (2)$$


where z_t are exogeneous variables that affect both interest rates and exchange rates. The nominal exchange rate s_t is defined as the units of Swiss franc per unit of a foreign currency. Thus, if s_t declines, the Swiss franc appreciates in nominal terms. The structural innovations (η_t, ε_t) are interpreted as an exchange rate and a monetary policy shock, respectively. The structural shocks are assumed to have a mean of zero and be uncorrelated with each other and with the exogeneous variables z_t . Our interest lies in identifying the parameter α in the exchange rate response function given by Equation (1).

Empirically, we only observe equilibria of exchange rates and interest rates simultaneously, which makes it impossible to identify the response function of the exchange rates to interest rate changes with standard regression techniques. The naively observed relation will depend on the type of shock that hits the system. For example, a positive economic shock likely induces higher interest rates and a stronger Swiss franc, both due to the expected tightening of monetary policy. A negative risk shock will induce lower interest rates and a stronger Swiss franc, both due to the resulting safe-haven flows.

In fact, the rolling six-month correlation of the three-month CHF LIBOR in first differences and the EUR/CHF log return shown in Figure 1 is rather erratic due to the different type of shocks that affect both the exchange rate and the interest rate.

It is therefore important to adequately account for the feedback between interest rates and exchange rates. Using basic regression methods will result in biased estimates. To robustly infer the causal impact of policy rate changes on exchange rates, we employ the IH approach of Rigobon (2003) and Rigobon and Sack (2004). This approach is outlined in the next section.

Figure 1. Rolling Window Correlation for EUR/CHF and the Three-Month CHF LIBOR



Note: The chart shows the six-month rolling window correlation between Δ_s^{EURCHF} and $\Delta_t^{\text{3M-LIBOR}}$ from January 1, 2000 to August 31, 2011. The exchange rate is transformed by using the first difference of the log and the interest rate is transformed by the first difference.

3.2 Identification-through-Heteroskedasticity Methodology

Assume that the changes in the policy rate Δi_t and the exchange rate Δs_t are described by the above system of Equations (1) and (2), respectively. The simultaneous causality between Δi_t and Δs_t means that not all the parameters are exactly identified. The reduced form of the equation system, derived in Appendix B.1, has more unknown parameters than there are coefficients in the reduced form. Rigobon (2003) and Rigobon and Sack (2004) suggest a way to identify the exchange rate response α to a monetary policy announcement shock based on the heteroskedasticity of the shock. The idea is to look at the differences in the covariance structure of Δi_t and Δs_t for days with a monetary policy announcement (MPA) and days with no policy announcement.

We follow Rigobon and Sack (2004) in assuming that the monetary policy shock on days of MPAs of the SNB is relatively more important:

$$\begin{aligned}\sigma_{\varepsilon,P}^2 &> \sigma_{\varepsilon,\tilde{P}}^2, \\ \sigma_{\eta,P}^2 &= \sigma_{\eta,\tilde{P}}^2.\end{aligned}$$

Consequently, we split the data in two subsets that are assumed to differ only in terms of the variance of the monetary policy shock. Let P and $\tilde{P}$ denote the two subsamples. P is the set of days with an MPA, while $\tilde{P}$ contains non-MPA days. Following Rigobon and Sack (2004), we assign the day before the MPA to the subsample $\tilde{P}$ to reduce unpredictable influences from other economic shocks with respect to monetary policy actions. We assess the robustness of this choice in Appendix C.5. Since Switzerland is a small open economy and highly economically connected to Europe, monetary policy from other major economies affects Swiss exchange rates. We believe that focusing only on two days in the whole sample (non-policy sample with previous days' values and policy sample with MPA days) is therefore an advantage to isolate Swiss monetary policy surprises from other shocks (e.g., monetary policy surprises by the ECB). The key assumption is that the monetary policy innovations in the two sets have different variances, and the structural parameters and the variance of the exchange rate innovations are

unchanged.⁶ A monetary policy announcement by the SNB includes a monetary policy decision about its policy rate and a conditional inflation forecast. Furthermore, a forecast of Swiss real economic activity is provided. Since we identify monetary policy shocks based on monetary policy announcements, our identified shock reflects both components—the change of monetary policy and the change of the outlook of the Swiss economy.

Clearly, identification is complicated by the fact that interest rates and exchange rates are affected by a common set of exogenous variables z_t . An important benefit of the IH methodology is that potentially confounding variables need not be explicitly considered in order to identify α . As shown by Rigobon and Sack (2004), it is sufficient to impose the assumption that the variance of the exogenous shocks does not differ across the two subsamples:

$$\sigma_{z,P}^2 = \sigma_{z,\bar{P}}^2.$$

Given these assumptions, it can be shown that the difference between the covariance on MPA days and the covariance on non-MPA days is given by (see Appendix B.1 for a derivation):

$$\Delta\Omega = \lambda \begin{bmatrix} 1 & \alpha \\ \alpha & \alpha^2 \end{bmatrix} = \lambda \begin{bmatrix} 1 \\ \alpha \end{bmatrix} \begin{bmatrix} 1 & \alpha \end{bmatrix},$$

where

$$\lambda = \frac{\sigma_{\eta,P}^2 - \sigma_{\eta,\bar{P}}^2}{(1 - \alpha\beta)^2}. \quad (3)$$

That is, by examining the change in the covariance, we can isolate the policy impact parameter α and purge all other influences that are assumed to have an equal covariance structure on MPA and non-MPA days.

We can extend the model by introducing an additional exchange rate equation that allows us to analyze the policy effect on

⁶A recent paper by Jarocinski and Karadi (2020) deconstructs monetary policy surprises into information about the future stance of monetary policy and the economic outlook using high-frequency data. This is a relevant decomposition from which we have to abstract because we use daily data and the assumption about a switch in variances to identify monetary policy shocks.

EUR/CHF and USD/CHF in a single system. In this case, the change in the covariance is given by (see Appendix B.2 for a derivation):

$$\Delta\Omega = \lambda \begin{bmatrix} 1 & \alpha_1 & \alpha_2 \\ \alpha_1 & \alpha_1^2 & \alpha_1\alpha_2 \\ \alpha_2 & \alpha_1\alpha_2 & \alpha_2^2 \end{bmatrix} = \lambda \begin{bmatrix} 1 \\ \alpha_1 \\ \alpha_2 \end{bmatrix} \begin{bmatrix} 1 & \alpha_1 & \alpha_2 \end{bmatrix},$$

with

$$\lambda = \frac{\sigma_{\eta,P}^2 - \sigma_{\eta,\tilde{P}}^2}{(1 - \alpha_1\beta_1 - \alpha_2\beta_2)^2}.$$

The next section outlines our estimation approach.

3.3 Estimation Methods

Rigobon and Sack (2004) show that there are two ways to implement the estimation: using instrumental variables (IV) or using the generalized method of moments (GMM). In what follows, we focus on the GMM approach. To check robustness, we also implement the IV estimator, with consistent findings presented in Appendix C.

The GMM approach matches the model-implied variances and covariances with their empirical counterparts. For the single exchange rate case, we obtain three moment conditions in two unknown parameters (see Appendix B.1 for a derivation):

$$g(\lambda, \alpha) = \begin{bmatrix} \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta i_t^2 - \lambda \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta i_t \Delta s_t - \lambda \alpha \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta s_t^2 - \lambda \alpha^2 \end{bmatrix}.$$

In a system with two exchange rates, we obtain two additional terms for the variance of the second exchange rate and its covariance with the interest rate change. On top of that, we obtain another moment condition corresponding to the covariance between the two exchange rates. This results in six moment conditions in three unknown parameters (see Appendix B.2 for a derivation):

$$g(\lambda, \alpha_1, \alpha_2) = \begin{bmatrix} \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\bar{P}}} \delta_t^{\bar{P}} \right) \Delta i_t^2 - \lambda \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\bar{P}}} \delta_t^{\bar{P}} \right) \Delta i_t \Delta s_{1t} - \lambda \alpha_1 \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\bar{P}}} \delta_t^{\bar{P}} \right) \Delta s_{1t}^2 - \lambda \alpha_1^2 \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\bar{P}}} \delta_t^{\bar{P}} \right) \Delta i_t \Delta s_{2t} - \lambda \alpha_2 \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\bar{P}}} \delta_t^{\bar{P}} \right) \Delta s_{2t}^2 - \lambda \alpha_2^2 \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\bar{P}}} \delta_t^{\bar{P}} \right) \Delta s_{1t} s_{2t} - \lambda \alpha_1 \alpha_2 \end{bmatrix}.$$

The system with two exchange rates allows us to test for the equality of the impact coefficients α_1 and α_2 . Given that the null hypothesis of the equality of the effects cannot be rejected, we estimate a constrained model with $\alpha = \alpha_1 = \alpha_2$. Note that we expect equal effects on the EUR/CHF and USD/CHF exchange rate because it seems implausible that the EUR/USD exchange rate is affected by SNB policy rate decisions. In the constrained model, we have six moment conditions in two unknown parameters.

We estimate the unknown parameters by iterated efficient non-linear GMM, as detailed in Appendix B.3.

4. Data

Policy rates have been relatively stable internationally and in particular in Switzerland after having reached a lower bound in the aftermath of the global financial crisis. Therefore, we have to resort to historical monetary policy cycles spanning from January 1, 2000 to August 31, 2011 at daily frequency. The sample begins after the SNB changed its monetary policy concept in December 1999. After 25 years of monetary targeting, the SNB implemented a new monetary policy concept in December 1999. The new concept built on three elements. The first element was considering price stability to be compatible with annual CPI inflation of less than 2 percent. The second element was a conditional inflation forecast that is published on a quarterly basis. The forecast is based on the assumption that the policy rate set when the forecast is published will remain constant over the entire three-year forecast horizon. The third element


was an operational target range for the three-month CHF LIBOR.⁷ The SNB began to announce a target band for the three-month CHF LIBOR, which was chosen to be consistent with a medium-term inflation rate of below 2 percent. The sample ends prior to the EUR/CHF minimum exchange rate regime, which was introduced on September 6, 2011. This event was the beginning of a significant regime change of SNB's monetary policy, where besides negative policy rates FX interventions became the second pillar of SNB's monetary policy tools.⁸ In the sample, the SNB implemented its monetary policy mainly using one-week repurchase agreement operations to control the three-month CHF LIBOR. Typically, the SNB aimed to keep the reference rate in the middle of the target range. Figure 2 shows the SNB target range, the three-month CHF LIBOR, and the Swiss exchange rates in the left and right panels, respectively.

MPA Days. The IH methodology requires the sample to be split into MPA and non-MPA days to identify the slope of the exchange rate response function. Our time frame covers 56 monetary policy decisions made by the SNB. At 26 policy meetings, the target range of the three-month CHF LIBOR was changed. Monetary policy decisions are made on a quarterly basis by the SNB Governing Board at its monetary policy assessment. The monetary policy decision is announced in a press release. At the June and December monetary policy assessments, the members of the Governing Board also explain the monetary policy decision at a press conference. A monetary policy assessment consists of a monetary policy decision about the policy rate and a conditional inflation forecast. Furthermore, a forecast of real economic activity is provided. From 2003 until 2011, every regularly scheduled monetary policy meeting contains the conditional forecast for inflation over the next three years as well as the forecast for real economic activity. For June and December meetings, the detailed chart about conditional inflation

⁷Note that on June 13, 2019, the target range for the three-month CHF LIBOR was replaced by the SNB policy rate.

⁸The SNB announced that the minimum exchange rate is enforced with the utmost determination via buying foreign currency in unlimited quantities. The policy rate did not change during the minimum exchange rate regime and the volatility of the EUR/CHF exchange rate was low, which makes our identification approach not applicable in this time period.

Figure 2. SNB Target Range, Three-Month CHF LIBOR, and Swiss Franc Exchange Rates



Note: The chart shows the SNB target range for the three-month CHF LIBOR, the three-month CHF LIBOR and Swiss franc exchange rates on a daily basis from January 1, 2000 to August 31, 2011.

is presented in the press release “Introductory Remarks, News Conference,” whereas for March and September meetings the chart is presented in the press release “Monetary Policy Assessment.” The reason for this is that the June and December meetings have a press conference shortly after the official announcement, whereas March and September meetings do not have a press conference. Importantly, inflation and economic outlook forecasts are provided each meeting.⁹ Furthermore, 9 out of the 56 policy meetings were unscheduled special meetings. Table A.1 in Appendix A provides a detailed overview of the monetary policy announcement meetings.¹⁰

Policy Rate. We use the daily change in the three-month CHF LIBOR to measure the policy rate changes Δi_t , because it is the monetary policy instrument that is controlled by the SNB. While longer-term interest rates are also relevant for exchange rate dynamics, our analysis takes the perspective of a policymaker that is interested in the impact its monetary policy decision. Since the SNB targeted the three-month CHF LIBOR, the rate directly incorporates policy rate expectations between day t and $t + 90$. Hence, absent other shocks, the three-month CHF LIBOR will change only on MPA day t if the markets are surprised by the outcome of the SNB’s monetary policy decision.¹¹

⁹To check whether it is appropriate to combine MPAs with and without press conference, we analyzed the uncertainty around SNB policy meetings in high frequency, using the spread between the best bid and offer price of EUR/CHF and USD/CHF as a proxy for the arrival of new information. Typically, spreads increase during the arrival of new information, since uncertainty is large and market participants do not want to be on the wrong side of a transaction. As market participants interpret and digest the new information, spreads typically decrease to their normal levels. For the SNB policy meetings, we see a widening of the spread over the initial release, but not during the press conference (which is usually 30 minutes after the release). We conclude that the latter does not include substantial additional information and therefore we do not differentiate further between MPAs with and without a press conference. To compute the spread, we resample best bid and offer prices at a one-minute frequency over a window from 100 minutes before to 100 minutes after an announcement.

¹⁰The non-MPA sample contains two days where the ECB and the Fed held monetary policy announcement meetings (ECB: December 6, 2001; Fed: February 2, 2000). Excluding these days does not alter the results.

¹¹Note that for approximately 50 percent of the MPAs, we use the three-month CHF LIBOR change between t and $t + 1$ to measure unexpected policy rate changes. The reason for this is that the three-month CHF LIBOR is fixed at approximately 11:00 a.m. London time (12:00/13:00 Zurich daylight

Measuring the Expectations Component in Longer-Term Interest Rates. We decompose the Swiss government bond yield curve into two components: the expected average future short-term interest rates and the term premium. For this purpose, we use the term structure model and estimation procedure of Adrian, Crump, and Moench (2013). For each maturity, we calculate the expected component as the expected average future one-year interest rate. For decomposing the response of long-term expectations about future short-term interest rates, we use Swiss government bonds prices sampled at their end-of-day value.

Exchange Rates. As exchange rate variables we use both the EUR/CHF and USD/CHF exchange rate. As outlined above, these exchange rates measure the units of Swiss franc per unit of euro and U.S. dollar, respectively. Thus, if the exchange rate declines, the Swiss franc appreciates in nominal terms. The exchange rates are sampled at the same points in time as the interest rate variable. For the three-month CHF LIBOR, the relevant time of day is 11:00 a.m. London time, when the LIBOR is approximately fixed. For every day in our sample, we retrieve the best bid and offer prices on the interbank market at the respective point in London time. We then compute the daily exchange rate return as the difference in the logarithm of the mid-price. Table 1 provides an overview on the source and transformations of the interest rate and exchange rate variables.

Descriptive Statistics. The descriptive statistics for the subsamples P and $\tilde{P}$ are reported in Table 2. Note that the standard deviation of the three-month CHF LIBOR significantly increases by a factor larger than two on MPA days compared to the standard deviation on non-MPA days. The standard deviations of the CHF exchange rate returns increase as well, but the magnitude is much smaller. Thus, as expected, the descriptive statistics support the assumption that the variance of the monetary policy shock on MPA days relatively increases.

savings/standard time), whereas the monetary policy decisions were alternately announced before and after the LIBOR fixing. Between 2000 and 2010, the SNB policy announcements took place at 9:30 a.m. Zurich time in June and December (we use the three-month CHF LIBOR change $t - 1$ to t) and at 14:00 p.m. Zurich time in March and September (we use the three-month CHF LIBOR change t to $t + 1$). Unscheduled announcements were released at approximately 13:00 p.m. (we use the three-month CHF LIBOR change t to $t + 1$).

Table 1. Overview of the Financial Variables

Name	Description	Transformation	Unit
$\Delta_t^{3M-LIBOR}$	3-Month LIBOR Rate for CHF Deposits	First Difference	Ppts
Δ_s^{EURCHF}	EUR/CHF Exchange Rate	Log Return	Pct
Δ_s^{USDCHF}	USD/CHF Exchange Rate	Log Return	Pct
Δ_t^{2Y}	2-Year Swiss Government Bond Yield		Pct
Δ_t^{3Y}	3-Year Swiss Government Bond Yield		Pct
Δ_t^{5Y}	5-Year Swiss Government Bond Yield		Pct
Δ_t^{7Y}	7-Year Swiss Government Bond Yield		Pct
Δ_t^{10Y}	10-Year Swiss Government Bond Yield		Pct
Note: The table provides details of the variables used for estimation. The data source for the three-month CHF LIBOR and the Swiss government bond yields is Bloomberg L.P., whereas for the exchange rates we rely on EBS Data Mine data.			

Table 2. Descriptive Statistics

	Mean	Median	Minimum	Maximum	SD
<i>A. MPA Days</i>					
$\Delta_t^{3M-LIBOR}$	-0.037	-0.005	-0.612	0.282	0.122
Δ_s^{EURCHF}	0.135	0.095	-1.221	3.511	0.666
Δ_s^{USDCHF}	-0.012	-0.012	-1.833	2.629	0.914
<i>B. Non-MPA Days</i>					
$\Delta_t^{3M-LIBOR}$	-0.002	0.000	-0.135	0.070	0.029
Δ_s^{EURCHF}	-0.130	-0.051	-2.968	0.850	0.552
Δ_s^{USDCHF}	-0.267	-0.246	-4.684	1.191	0.892
Note: The table provides descriptive statistics for the exchange rates and the policy rate variable for both subsamples. The sample size for each subsample is 56, according to the number of monetary policy announcements by the SNB from January 1, 2000 to August 31, 2011.					

5. Results

5.1 *Exchange Rate Response to a Monetary Policy Shock*

This section discusses the main result of the present analysis, namely, the estimated response of exchange rates to a monetary policy shock (i.e., parameter α in Equation (1)).

Exchange Rate Response. Table 3 shows the IH-GMM-estimation results. Our baseline specifications use the three-month CHF LIBOR as the policy variable. The two left-most columns show the results of estimating two separate equations for EUR/CHF and USD/CHF. The significance tests and standard deviations are asymptotic. In response to a 100 bp increase in the three-month LIBOR, the point estimates suggest that the Swiss franc appreciates by 2.1 percent and 1.9 percent on the same day against the euro and the U.S. dollar, respectively.¹² The third and fourth columns show that consistent estimates result when estimating the joint system that includes both EUR/CHF and USD/CHF. While the point estimates are fairly similar, using all available information reduces the standard error of the estimation.

Next, we test the null hypothesis that the effects are the same for both the EUR/CHF and USD/CHF exchange rates. The results for the corresponding Wald test are shown in Table 3. The null hypothesis regarding the equality of the effects cannot be rejected. This result is sensible from an economic perspective because differing effects would imply that the monetary policy announcements made by the SNB are able to move the EUR/USD exchange rate.

As the null hypothesis of equality cannot be rejected, we restrict the impact coefficients in the EUR/CHF and USD/CHF equations to be equal. The results shown in the last column of Table 3 indicate that this further improves the precision of the estimates. The point estimate for α suggests that in response to a 100 bp increase in the three-month CHF LIBOR, the Swiss franc appreciates by 2.0

¹²Table A.1 in Appendix A reports the change of the three-month CHF LIBOR (in bp) for each monetary policy announcement day. For example, on November 6, 2008, the SNB decided to decrease the lower as well as the upper bound of the target range by 50 bp. The three-month CHF LIBOR decreased by 25 bp. On average, the absolute change of the three-month CHF LIBOR of monetary policy announcements dates is 7 bp.

Table 3. IH-GMM Estimates of the Exchange Rate Response

	Separate Equations		Joint Estimation		Restricted $\alpha_1 = \alpha_2$
	EUR/CHF	USD/CHF	EUR/CHF	USD/CHF	EUR/CHF USD/CHF
$\hat{\alpha}$	-2.1*** (0.46)	-1.9* (1.0)	-2.0*** (0.32)	-1.8** (1.0)	-2.0*** (0.3)
$\hat{\lambda}$	0.015** (0.0072)	0.015** (0.0072)	0.015** (0.0071)		0.015*** (0.0048)
Wald $H_0 : \alpha_1 = \alpha_2$ p-value			0.038 0.84		
J-statistic	0.066	0.032	0.18		0.22
p-value	0.8	0.86	0.98		0.99
No. Moment Conditions	3	3	6		6
No. Parameters	2	2	3		2
Note: The table shows the IH-GMM estimates. The policy rate variable is the three-month CHF LIBOR. The J-statistic reports the results of testing the null hypothesis regarding the validity of overidentifying restrictions. *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively. The numbers in parentheses are standard deviations.					

percent against the euro and the U.S. dollar, with the 95 percent confidence interval spanning from 1.4 percent to 2.5 percent.

The GMM table includes the estimates of the parameter λ , which represents the change in the variance of ϵ_t between MPA and non-MPA dates divided by a determinant. The IH approach works only if a change in the variance is present. The estimates indicate that this is indeed the case for all the specifications.


The remaining rows of Table 3 show the tests for the validity of the overidentifying restrictions. The null hypothesis regarding the validity of the restrictions cannot be rejected for any of the different specifications. Finally, the last row shows the number of iterations needed for the convergence of iterated efficient GMM. The results show that convergence is achieved fairly quickly.

The magnitude of our benchmark result for the exchange rate response lies within the range of effects reported in the literature and are similar to the findings of Grisse (2020) and Kugler (2020).

Relevance of Accounting for Simultaneity. Accounting for simultaneity is important. Using simple regression methods that do not disentangle the feedback between interest rates and exchange rates results in biased estimates. Regressing exchange rate returns on interest rate changes yields estimates significantly different from the IH estimates. Appendix C.1 presents the biased ordinary least squares (OLS) results. The OLS estimate for EUR/CHF is not significantly different from zero, while the USD/CHF coefficient has the opposite sign. We conclude that with standard regressions (or event studies at the daily level), one cannot correctly identify the causal impact of interest rates on exchange rates. As outlined in Section 3.1, this occurs because interest rates and foreign exchange rates react to each other during the day.



Robustness Checks. Table 4 and Appendix C include four robustness checks: First, we use the IH-IV estimator instead of the GMM (see Appendix C.2). Second, we sample exchange rates at the end of the day (5:00 p.m. New York time) rather than at the points in time when the three-month CHF LIBOR is fixed (11:00 a.m. London time) (see Appendix C.3). Third, we use estimates of the three-month constant maturity rate inferred from three-month CHF LIBOR futures as the policy rate (see Appendix C.4). The ICE LIBOR (intercontinental exchange London interbank offered rate) futures data as well as exchange rates are sampled at the close of the

Table 4. Estimates of $\hat{\alpha}$ from Alternative Specifications

Model	Policy	Rate	Samp.		$\hat{\alpha}$	SE	95% CI
Rest	LIBOR	CHF	Sync		-2.0	0.3	(-2.5 to -1.4)
Rest	LIBOR	CHF	End		-2.6	0.8	(-4.2 to -0.9)
Joint	LIBOR	EUR/CHF	Sync		-2.0	0.3	(-2.6 to -1.3)
Joint	LIBOR	EUR/CHF	End		-2.2	0.8	(-3.7 to -0.7)
Joint	LIBOR	USD/CHF	Sync		-1.8	1.0	(-3.7 to +0.2)
Joint	LIBOR	USD/CHF	End		-3.2	1.2	(-5.5 to -0.9)
Sep	LIBOR	EUR/CHF	Sync		-2.1	0.5	(-2.9 to -1.2)
Sep	LIBOR	EUR/CHF	End		-1.8	0.8	(-3.4 to -0.3)
Sep	LIBOR	USD/CHF	Sync		-1.9	1.0	(-3.9 to +0.2)
Sep	LIBOR	USD/CHF	End		-3.2	1.3	(-5.7 to -0.7)
Rest	Future	CHF	Sync		-3.1	1.3	(-5.7 to -0.4)
Joint	Future	EUR/CHF	Sync		-1.7	0.9	(-3.5 to +0.0)
Joint	Future	USD/CHF	Sync		-3.7	1.3	(-6.2 to -1.1)
Sep	Future	EUR/CHF	Sync		-1.5	0.9	(-3.2 to +0.2)
Sep	Future	USD/CHF	Sync		-3.7	1.3	(-6.2 to -1.2)
IV	LIBOR	EUR/CHF	Sync		-2.2	0.7	(-3.4 to -0.9)
IV	LIBOR	EUR/CHF	End		-1.9	0.7	(-3.3 to -0.5)
IV	LIBOR	USD/CHF	Sync		-2.4	1.0	(-4.3 to -0.5)
IV	LIBOR	USD/CHF	End		-2.7	1.0	(-4.6 to -0.7)
IV	Future	EUR/CHF	Sync		-1.2	0.7	(-2.6 to +0.1)
IV	Future	EUR/CHF	End		-1.8	0.8	(-3.3 to -0.3)
IV	Future	USD/CHF	Sync		-3.0	1.0	(-4.9 to -1.0)
IV	Future	USD/CHF	End		-3.3	1.1	(-5.4 to -1.2)

Note: “Rest”: restricted GMM. “Joint”: joint GMM. “Sep”: separate GMM. “Samp.”: exchange rate sampling. “Sync”: synchronized to policy variable (i.e., 11:00 a.m. London for LIBOR; 6:00 p.m. London for futures). “End”: closing rate (i.e., 5:00 p.m. New York).

Table 5. Term Structure Response

Parameter	$\hat{\alpha}$	SE		95% CI
2Y	0.34	0.11		(+0.13 to +0.55)
3Y	0.29	0.10		(+0.09 to +0.48)
5Y	0.22	0.09		(+0.42 to +0.40)
7Y	0.19	0.75		(+0.41 to +0.33)
10Y	0.15	0.07		(+0.01 to +0.28)
Note: IH-GMM estimation of the response of the expected average short-term interest rate to an unexpected policy rate change by 100 bp. Joint GMM estimation with 21 moment conditions for six parameters. $\hat{\lambda} = 0.009^{**}(0.0045)$. J-statistic for validity of overidentifying restrictions: 2.41 with a p-value of 0.99. The expected average short-term interest rates are generated variables and therefore the confidence intervals may be biased.				

day (6:00 p.m. London time). Fourth, we investigated the robustness of our results with respect to the choice of the non-policy sample.

Table 4 visually shows that our benchmark results (in the first row) are supported by alternative specifications. The confidence intervals can become larger for a few alternative specifications consistent with even more pronounced effects than our benchmark results suggest. For 18 out of 23 specifications, the point estimates are significantly different from zero at the 5 percent significance level. At the 10 percent level, all point estimates are significantly different from zero. We conclude that our key result that an unexpected increase in the policy rate leads to an appreciation of the Swiss franc is very robust.

5.2 Yield Curve Response to a Monetary Policy Shock

In this section, we analyze the impact of a monetary policy shock on the yield curve of Swiss government bonds. The expected average future 1-year interest rate is computed for 2-, 3-, 5-, 7-, and 10-year terms using the approach of Adrian, Crump, and Moench (2013). The interest rate expectation component is the financial market’s best forecast of short-term yields over the lifetime of the bond. This analysis will shed light on how an unexpected change in the three-month CHF LIBOR affects the expectation of the expected average future short-term interest rates.

Table 5 shows the IH-GMM joint model estimates of the response of the expected average future short-term interest rate to a monetary

policy shock that changes the policy rate by 100 bp. The estimates reveal that expected average future short-term rates are affected by monetary policy shocks. For the two-year maturity, we find that a 100 bp increase in the three-month CHF LIBOR increases the expected average short-term rate by 34 bp. The effect declines to 15 bp for the 10-year-ahead expected short-term rate.¹³ The findings are in line with Grisse and Schumacher (2018), who also report that short-term and longer-term interest rates tend to move in the same direction but not one-for-one.

We conclude that the identified monetary policy shock significantly affects medium- to longer-term policy rate expectations, which in turn affects the exchange rate.

6. Conclusions

This paper identifies the short-term impact of monetary policy announcements made by the SNB on the Swiss franc exchange rates and on the Swiss yield curve. Our results robustly show that a monetary policy announcement that increases the policy rate appreciates the nominal Swiss franc on the same day. The null hypothesis that monetary policy announcement shocks do not affect the exchange value of the Swiss franc is clearly rejected.

Importantly, we also show that simple methods that do not adequately account for the simultaneous relation of exchange rates and interest rates yield biased and typically non-significant results. This may lead to the incorrect conclusion that monetary policy changes do not affect the Swiss franc.

Moreover, we find that our identified monetary policy announcement shock affects the expected path of future short-term interest rates, which in turn influence the valuation of the Swiss franc.

¹³We also estimated the impact of a monetary policy shock on the term premium at different maturities. We find that an increase in the policy rate typically leads to a small effect (0.05 percentage point for a 100 bp increase) in the term premium for short- and medium-term maturities. For longer maturities, this effect decreases. The pattern of small increases at shorter- to medium-term and the decrease for longer-term maturities is in line with the results of Soederlind (2010).

The magnitude of the benchmark results lies within the range of estimates in the literature. In this study we use data on two conventional monetary policy cycles in the historical sample from 2000 to 2011. However, recent contributions have suggested that in an environment of low interest rates markets might have become more sensitive to restrictive monetary policy shifts. It would be interesting to analyze the effects on the Swiss franc and longer-term interest rates resulting from changes in the monetary policy stances of other central banks such as the ECB or Fed. All these possible extensions are, however, left for future research.

Table A.1 provides details of the MPA days of the SNB from January 1, 2000 to August 31, 2011.

Table A.1. SNB Monetary Policy Announcement Dates

Year	Day	Comment	Δi	Δs^E	Δs^U
2000	20 Jan	R: Unchanged	3.3	-0.17	-0.53
	03 Feb	S: $\Delta \underline{B} = 50$ (1.75), $\Delta \overline{B} = 50$ (2.75)	14.7	-0.25	-1.63
	23 Mar	R: $\Delta \underline{B} = 75$ (2.50), $\Delta \overline{B} = 75$ (3.50)	28.2	-0.60	-1.69
	15 Jun	R: $\Delta \underline{B} = 50$ (3.00), $\Delta \overline{B} = 50$ (4.00)	1.3	-0.20	0.08
	14 Sep	R: Unchanged	-4.8	0.71	0.15
2001	08 Dec	R: Unchanged	-2.0	-0.04	0.06
	22 Mar	R: $\Delta \underline{B} = -25$ (2.75), $\Delta \overline{B} = -25$ (3.75)	-11.2	-0.11	0.74
	14 Jun	R: Unchanged	13.2	-0.10	-0.95
	17 Sep	R: $\Delta \underline{B} = -50$ (2.25), $\Delta \overline{B} = -50$ (3.25)	-15.7	-2.16	-5.11
	24 Sep	S: $\Delta \underline{B} = -50$ (1.75), $\Delta \overline{B} = -50$ (2.75)	-25.2	1.29	1.15
2002	07 Dec	R: $\Delta \underline{B} = -50$ (1.25), $\Delta \overline{B} = -50$ (2.25)	-7.3	0.09	0.64
	21 Mar	R: Unchanged	-0.7	-0.24	0.17
	02 May	S: $\Delta \underline{B} = -50$ (0.75), $\Delta \overline{B} = -50$ (1.75)	-19.8	0.07	0.42
	14 Jun	R: Unchanged	-0.8	0.13	-0.16
	26 Jul	S: $\Delta \underline{B} = -50$ (0.25), $\Delta \overline{B} = -50$ (1.25)	-19.2	-0.03	1.76
2003	19 Sep	R: Unchanged	0.0	-0.40	-1.32
	13 Dec	R: Unchanged	2.7	-0.01	-0.53
	06 Mar	S: $\Delta \underline{B} = -25$ (0.00), $\Delta \overline{B} = -50$ (0.75)	-21.5	0.40	0.41
	20 Mar	R: Unchanged	-0.4	0.00	-0.45
	13 Jun	R: Unchanged	-0.7	-0.25	-1.15
2004	18 Sep	R: Unchanged	0.0	0.13	0.39
	12 Dec	R: Unchanged	-0.3	0.02	-0.40
	18 Mar	R: Unchanged	-0.2	-0.32	-1.55
	17 Jun	R: Unchanged $\Delta \underline{B}$, $\Delta \overline{B} = 25$ (1.00)	6.2	-0.76	-1.13
	16 Sep	R: $\Delta \underline{B} = 25$ (0.25), $\Delta \overline{B} = 25$ (1.25)	-0.8	0.17	-0.09
2005	16 Dec	R: Unchanged	-3.2	0.29	1.43
	17 Mar	R: Unchanged	-0.8	0.08	0.36

(continued)

Table A.1. (Continued)

Year	Day	Comment	Δi	Δs^E	Δs^U
2006	16 Jun	R: Unchanged	0.0	0.01	0.09
	15 Sep	R: Unchanged	-0.7	0.27	0.78
	15 Dec	R: $\Delta \bar{B} = 25$ (0.50), $\Delta \bar{B} = 25$ (1.50)	-6.7	0.37	0.55
	16 Mar	R: $\Delta \bar{B} = 25$ (0.75), $\Delta \bar{B} = 25$ (1.75)	-1.4	0.46	-0.47
	15 Jun	R: $\Delta \bar{B} = 25$ (1.00), $\Delta \bar{B} = 25$ (2.00)	0.5	0.26	-0.07
2007	14 Sep	R: $\Delta \bar{B} = 25$ (1.25), $\Delta \bar{B} = 25$ (2.25)	2.7	0.32	0.04
	14 Dec	R: $\Delta \bar{B} = 25$ (1.50), $\Delta \bar{B} = 25$ (2.50)	1.0	0.17	0.70
	15 Mar	R: $\Delta \bar{B} = 25$ (1.75), $\Delta \bar{B} = 25$ (2.75)	-0.5	0.19	0.07
	14 Jun	R: $\Delta \bar{B} = 25$ (2.00), $\Delta \bar{B} = 25$ (3.00)	0.1	0.15	0.13
	13 Sep	R: $\Delta \bar{B} = 25$ (2.25), $\Delta \bar{B} = 25$ (3.25)	-9.2	-0.05	0.09
2008	13 Dec	R: Unchanged	1.3	0.07	0.57
	13 Mar	R: Unchanged	-1.5	0.05	-0.53
	19 Jun	R: Unchanged	-6.4	0.69	0.90
	18 Sep	R: Unchanged	1.3	0.30	0.15
	08 Oct	S: $\Delta \bar{B} = -25$ (2.00), $\Delta \bar{B} = -25$ (3.00)	1.3	-0.68	-1.14
2009	06 Nov	S: $\Delta \bar{B} = -50$ (1.50), $\Delta \bar{B} = -50$ (2.50)	-25.3	-0.14	1.71
	20 Nov	S: $\Delta \bar{B} = -100$ (0.50), $\Delta \bar{B} = -100$ (1.50)	-61.2	0.59	0.89
	11 Dec	S: $\Delta \bar{B} = -50$ (0.00), $\Delta \bar{B} = -50$ (1.00)	-27.3	1.28	-1.25
	12 Mar	R: Unchanged $\Delta \bar{B}$, $\Delta \bar{B} = -25$ (0.75)	-3.8	3.30	2.75
	18 Jun	R: Unchanged	0.0	0.28	0.60
2010	17 Sep	R: Unchanged	-0.3	-0.13	-0.38
	10 Dec	R: Unchanged	0.0	-0.03	-0.07
	11 Mar	R: Unchanged	0.0	0.04	-0.16
	17 Jun	R: Unchanged	0.7	-1.07	-1.70
	16 Sep	R: Unchanged	-0.7	1.76	1.23
2011	16 Dec	R: Unchanged	-0.1	-0.15	-0.38
	17 Mar	R: Unchanged	0.0	-0.17	-1.07
	16 Jun	R: Unchanged	0.0	-0.43	-0.64
	03 Aug	S: Unchanged $\Delta \bar{B}$, $\Delta \bar{B} = -50$ (0.25)	-3.7	1.89	1.07
Note: The table provides details of the MPA days of the SNB from January 1, 2000 to August 31, 2011. $\Delta \bar{B}$ ($\Delta \bar{B}$) indicates the change in lower (upper) bound of the three-month CHF LIBOR target range (in bp). The level is in parentheses (in pct). R (S) stands for regular (special) MPAs. Δi , Δs^E , and Δs^U are the change in the three-month CHF LIBOR (in bp) and the return of the EUR/CHF and USD/CHF exchange rate (in pct) on the MPA day, respectively.					

Appendix B. Estimation by GMM

This section details the estimation of the interest rate effect by the generalized method of moments. Section B.1 considers the case of a single exchange rate equation and derives the model-implied change in the covariance matrix between monetary policy announcement days and non-monetary-policy announcement days. From this, we can deduce the moment conditions that enables us to estimate the policy effect. Section B.2 extends the system by adding an additional exchange rate equation and works out the corresponding moment conditions.

B.1 Single Exchange Rate Equation

Let us start with the equation system given by (1) and (2). In matrix form this can be represented by

$$\begin{bmatrix} 1 & -\beta \\ -\alpha & 1 \end{bmatrix} \begin{bmatrix} \Delta i_t \\ \Delta s_t \end{bmatrix} = \begin{bmatrix} \gamma_i \\ \gamma_s \end{bmatrix} z_t + \begin{bmatrix} \epsilon_t \\ \eta_t \end{bmatrix}.$$

The reduced form of the system is

$$\begin{aligned} \begin{bmatrix} \Delta i_t \\ \Delta s_t \end{bmatrix} &= \frac{1}{(1 - \alpha\beta)^2} \begin{bmatrix} 1 & \beta \\ \alpha & 1 \end{bmatrix} \left(\begin{bmatrix} \gamma_i \\ \gamma_s \end{bmatrix} z_t + \begin{bmatrix} \epsilon_t \\ \eta_t \end{bmatrix} \right) \\ &= \frac{1}{(1 - \alpha\beta)^2} \begin{bmatrix} (\gamma_i + \beta\gamma_s)z_t + \epsilon_t + \beta\eta_t \\ (\alpha\gamma_i + \gamma_s)z_t + \alpha\epsilon_t + \eta_t \end{bmatrix}. \end{aligned}$$

The covariance is given by

$$\Omega = \begin{bmatrix} \sigma_{\Delta i}^2 & \sigma_{\Delta i \Delta s} \\ \sigma_{\Delta i \Delta s} & \sigma_{\Delta s}^2 \end{bmatrix},$$

where

$$\begin{aligned} \sigma_{\Delta i}^2 &= \frac{1}{(1 - \alpha\beta)^2} [(\gamma_i + \beta\gamma_s)^2 \sigma_z^2 + \sigma_\epsilon^2 + \beta^2 \sigma_\eta^2] \\ \sigma_{\Delta i \Delta s} &= \frac{1}{(1 - \alpha\beta)^2} [(\gamma_i + \beta\gamma_s)(\alpha\gamma_i + \gamma_s) \sigma_z^2 + \alpha \sigma_\epsilon^2 + \beta \sigma_\eta^2] \\ \sigma_{\Delta s}^2 &= \frac{1}{(1 - \alpha\beta)^2} [(\gamma_i + \beta\gamma_s)^2 \sigma_z^2 + \alpha^2 \sigma_\epsilon^2 + \sigma_\eta^2]. \end{aligned}$$

When calculating the difference between the covariance of policy dates and the covariance of non-policy dates, the terms with σ_z^2 and σ_η^2 will cancel out, because their variance is assumed to be the same on policy and non-policy dates. The covariance difference implied by the model reduces to (see also Rigobon and Sack 2004, Equation 9):

$$\Delta\Omega = \lambda \begin{bmatrix} 1 & \alpha \\ \alpha & \alpha^2 \end{bmatrix} = \lambda \begin{bmatrix} 1 \\ \alpha \end{bmatrix} \begin{bmatrix} 1 & \alpha \end{bmatrix}, \quad (\text{B.1})$$

with

$$\lambda = \frac{\sigma_{\eta,P}^2 - \sigma_{\eta,\tilde{P}}^2}{(1 - \alpha\beta)^2}.$$

The empirical equivalent to this covariance matrix difference is

$$\Delta\hat{\Omega} = \hat{\Omega}_P - \hat{\Omega}_{\tilde{P}},$$

with

$$\begin{aligned} \hat{\Omega}_P &= \frac{1}{T_P} \sum_{t=1}^T \delta_t^P \Delta x_t \Delta x'_t, \\ \hat{\Omega}_{\tilde{P}} &= \frac{1}{T_{\tilde{P}}} \sum_{t=1}^T \delta_t^{\tilde{P}} \Delta x_t \Delta x'_t, \end{aligned}$$

where

$$\Delta x_t = \begin{bmatrix} \Delta i_t \\ \Delta s_t \end{bmatrix}$$

and δ_t^P and $\delta_t^{\tilde{P}}$ are dummies for policy and non-policy dates, respectively. From this,

$$\begin{aligned} \Delta\hat{\Omega} &= \hat{\Omega}_P - \hat{\Omega}_{\tilde{P}} & (\text{B.2}) \\ &= \frac{1}{T_P} \sum_{t=1}^T \delta_t^P \Delta x_t \Delta x'_t - \frac{1}{T_{\tilde{P}}} \sum_{t=1}^T \delta_t^{\tilde{P}} \Delta x_t \Delta x'_t \\ &= \sum_{t=1}^T \frac{1}{T_P} \delta_t^P \Delta x_t \Delta x'_t - \sum_{t=1}^T \frac{1}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \Delta x_t \Delta x'_t \end{aligned}$$

$$\begin{aligned}
&= \sum_{t=1}^T \left[\frac{1}{T_P} \delta_t^P \Delta x_t \Delta x'_t - \frac{1}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \Delta x_t \Delta x'_t \right] \\
&= \sum_{t=1}^T \left[\left(\frac{1}{T_P} \delta_t^P - \frac{1}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta x_t \Delta x'_t \right]. \tag{B.3}
\end{aligned}$$

We find the moment conditions by matching the variances and covariances in (B.1) with their empirical counterparts in (B.3). This result provides three moment conditions for the two unknown parameters, corresponding to the two variance terms and the covariance term:

$$g(\lambda, \alpha) = \begin{bmatrix} \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta i_t^2 - \lambda \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta i_t \Delta s_t - \lambda \alpha \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta s_t^2 - \lambda \alpha^2 \end{bmatrix}.$$

Identification requires that

$$E[g(\lambda, \alpha)] \neq 0 \quad \text{for} \quad \begin{bmatrix} \lambda \\ \alpha \end{bmatrix} \neq 0,$$

and

$$G = E \left[\frac{\partial g(\lambda, \alpha)}{\partial [\lambda \ \alpha]} \right] = \begin{bmatrix} -1 & 0 \\ -\alpha & -\lambda \\ -\alpha^2 & -2\lambda\alpha \end{bmatrix}$$

has full column rank 3. This requires $\lambda \neq 0$.

B.2 Two Exchange Rate Equations

This section augments the model by introducing another exchange rate equation, which allows us to analyze both the EUR/CHF and USD/CHF effects. For brevity, we disregard the exogenous variables. As in the single exchange rate equation case, these variables will drop out when computing the difference in the covariance between policy and non-policy dates.

The equation system with two exchange rate equations is given by

$$\begin{aligned}\Delta i_t &= \beta_1 \Delta s_{1t} + \beta_2 \Delta s_{2t} + \epsilon_t \\ \Delta s_{1t} &= \alpha_1 \Delta i_t + \eta_{1t} \\ \Delta s_{2t} &= \alpha_2 \Delta i_t + \eta_{2t},\end{aligned}$$

where Δi_t is the first difference in the interest rate, Δs_{1t} is the difference in the logarithm of the EUR/CHF exchange rate, and Δs_{2t} is the difference in the logarithm of the USD/CHF exchange rate. In matrix form this system can be represented as

$$\begin{bmatrix} 1 & -\beta_1 & -\beta_2 \\ -\alpha_1 & 1 & 0 \\ -\alpha_2 & 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta i_t \\ \Delta s_{1t} \\ \Delta s_{2t} \end{bmatrix} = \begin{bmatrix} \epsilon_t \\ \eta_{1t} \\ \eta_{2t} \end{bmatrix}.$$

The reduced form is

$$\begin{bmatrix} \Delta i_t \\ \Delta s_{1t} \\ \Delta s_{2t} \end{bmatrix} = \frac{1}{1 - \alpha_1 \beta_1 - \alpha_2 \beta_2} \begin{bmatrix} 1 & \beta_1 & \beta_2 \\ \alpha_1 & 1 - \alpha_2 \beta_2 & \alpha_1 \beta_2 \\ \alpha_2 & \alpha_2 \beta_1 & 1 - \alpha_1 \beta_1 \end{bmatrix} \begin{bmatrix} \epsilon_t \\ \eta_{1t} \\ \eta_{2t} \end{bmatrix}.$$

The individual equations are thus

$$\begin{aligned}\Delta i_t &= \frac{1}{1 - \alpha_1 \beta_1 - \alpha_2 \beta_2} [\epsilon_t - \beta_1 \eta_{1t} - \beta_2 \eta_{2t}] \\ \Delta s_{1t} &= \frac{1}{1 - \alpha_1 \beta_1 - \alpha_2 \beta_2} [\alpha_1 \epsilon_t + (1 - \alpha_2 \beta_2) \eta_{1t} + \alpha_1 \beta_2 \eta_{2t}] \\ \Delta s_{2t} &= \frac{1}{1 - \alpha_1 \beta_1 - \alpha_2 \beta_2} [\alpha_2 \epsilon_t + \alpha_2 \beta_1 \eta_{2t} + (1 - \alpha_2 \beta_1) \eta_{2t}].\end{aligned}$$

From these we can compute the three variances and three covariances of the variables. As most terms will again drop out when building the difference between policy dates and non-policy dates, we omit the formulas. The difference in the covariance is

$$\Delta \Omega = \lambda \begin{bmatrix} 1 & \alpha_1 & \alpha_2 \\ \alpha_1 & \alpha_1^2 & \alpha_1 \alpha_2 \\ \alpha_2 & \alpha_1 \alpha_2 & \alpha_2^2 \end{bmatrix} = \lambda \begin{bmatrix} 1 \\ \alpha_1 \\ \alpha_2 \end{bmatrix} \begin{bmatrix} 1 & \alpha_1 & \alpha_2 \end{bmatrix},$$

with

$$\lambda = \frac{\sigma_{\eta,P}^2 - \sigma_{\eta,\tilde{P}}^2}{(1 - \alpha_1\beta_1 - \alpha_2\beta_2)^2}.$$

The empirical equivalent to this covariance matrix difference is

$$\Delta\hat{\Omega} = \hat{\Omega}_P - \hat{\Omega}_{\tilde{P}},$$

with

$$\begin{aligned}\hat{\Omega}_P &= \frac{1}{T_P} \sum_{t=1}^T \delta_t^P \Delta x_t \Delta x_t', \\ \hat{\Omega}_{\tilde{P}} &= \frac{1}{T_{\tilde{P}}} \sum_{t=1}^T \delta_t^{\tilde{P}} \Delta x_t \Delta x_t',\end{aligned}$$

where

$$\Delta x_t = \begin{bmatrix} \Delta i_t \\ \Delta s_{1t} \\ \Delta s_{2t} \end{bmatrix}.$$

Again, matching the model-implied variances and covariance with their empirical counterparts gives six moment conditions in three unknown parameters:

$$g(\lambda, \alpha_1, \alpha_2) = \begin{bmatrix} \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta i_t^2 - \lambda \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta i_t \Delta s_{1t} - \lambda \alpha_1 \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta s_{1t}^2 - \lambda \alpha_1^2 \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta i_t \Delta s_{2t} - \lambda \alpha_2 \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta s_{2t}^2 - \lambda \alpha_2^2 \\ \left(\frac{T}{T_P} \delta_t^P - \frac{T}{T_{\tilde{P}}} \delta_t^{\tilde{P}} \right) \Delta s_{1t} \Delta s_{2t} - \lambda \alpha_1 \alpha_2 \end{bmatrix}.$$

We now have six equations in three unknown parameters. Identification requires that

$$E[g(\lambda, \alpha_1, \alpha_2)] \neq 0 \quad \text{for} \quad \begin{bmatrix} \lambda \\ \alpha_1 \\ \alpha_2 \end{bmatrix} \neq 0,$$

and

$$G = E \left[\frac{\partial g(\lambda, \alpha_1, \alpha_2)}{\partial [\lambda \ \alpha_1 \ \alpha_2]} \right] = \begin{bmatrix} -1 & 0 & 0 \\ -\alpha_1 & -\lambda & 0 \\ -\alpha_1^2 & -2\lambda\alpha_1 & 0 \\ -\alpha_2 & 0 & -\lambda \\ -\alpha_2^2 & 0 & -2\lambda\alpha_2 \\ -\alpha_1\alpha_2 & -\lambda\alpha_2 & -\lambda\alpha_1 \end{bmatrix}$$

has full column rank 3. This requires $\lambda \neq 0$.

B.3 Iterated Efficient GMM

Given the moment conditions, we can estimate the parameters by iterated efficient GMM.¹⁴ The sample moment condition for arbitrary parameters is

$$g_T(\lambda, \alpha) = \frac{1}{T} \sum_{t=1}^T g(\Delta i_t, \Delta s_t; \lambda, \alpha).$$

The efficient GMM estimator is given by

$$\hat{\theta}(\hat{W}) = \arg \min J(\lambda, \alpha, \hat{W}) = T g_T(\lambda, \alpha)' \hat{W} g_T(\lambda, \alpha),$$

where $\theta = [\lambda \ \alpha]'$, $\hat{W} = \hat{S}^{-1}$, such that $\hat{S} \xrightarrow{p} S = \text{avar}(g_T(\lambda, \alpha))$.

Given the consistent estimates $\hat{\lambda}$ and $\hat{\alpha}$ of λ and α , respectively, a heteroskedasticity (HC) estimate of S is

$$\hat{S}_{HC} = \frac{1}{T} \sum_{t=1}^T g(\hat{\lambda}, \hat{\alpha}) g(\hat{\lambda}, \hat{\alpha})'.$$

¹⁴The notations in this section loosely follow Zivot and Wang (2007, Section 21.6).

For the efficient GMM estimator, we use $\hat{W} = \hat{S}_{HC}^{-1}$, and it can be shown that

$$\begin{aligned}\hat{\theta}(\hat{S}_{HC}^{-1}) &\xrightarrow{p} \theta \\ \text{avar}(\hat{\theta}(\hat{S}_{HC}^{-1})) &= (\hat{G}' \hat{S}_{HC}^{-1} \hat{G})^{-1} \\ \hat{G} &= \frac{1}{T} \sum_{t=1}^T \frac{\partial g(\hat{\theta}(\hat{W}))}{\partial \theta'}.\end{aligned}$$

We use the iterated efficient GMM, stopping once the change in the moment norm is smaller than 1e-12, which is usually achieved in a few iterations.

The J-statistic for the validity of the moment conditions is given by

$$J = T g_T(\hat{\theta}(\hat{S}_{HC}^{-1}))' \hat{S}_{HC}^{-1} g_T(\hat{\theta}(\hat{S}_{HC}^{-1}))$$

and has a χ^2 -distribution with one degree of freedom.

Appendix C. Robustness

C.1 Results with OLS Estimator

We regress the exchange rate returns on three-month CHF LIBOR rate changes. The results are shown in Table C.2. The estimation results suggest that the estimates significantly differ from the IH estimates for EUR/CHF (point estimate: -0.4 , standard error: 0.3) and USD/CHF (point estimate: 2.5 , standard error: 0.5). The OLS estimates are not significantly different from zero, except for the estimates for USD/CHF. For USD/CHF, however, the OLS estimate is positive.

C.2 Results with Instrumental-Variable Estimator

Estimation by instrumental variables following Rigobon and Sack (2004) is based on using the instruments (w_i, w_s) for the endogenous variables $(\Delta i, \Delta s)$

$$w_i \equiv \begin{bmatrix} \Delta i_P \\ -\Delta i_{\bar{P}} \end{bmatrix}, \quad \Delta i \equiv \begin{bmatrix} \Delta i_P \\ \Delta i_{\bar{P}} \end{bmatrix}, \quad w_s \equiv \begin{bmatrix} \Delta s_P \\ -\Delta s_{\bar{P}} \end{bmatrix}, \quad \Delta s \equiv \begin{bmatrix} \Delta s_P \\ \Delta s_{\bar{P}} \end{bmatrix}. \quad (\text{C.1})$$

Table C.1. GMM Estimates with End-of-Day Exchange Rates

	Separate Equations		Joint Estimation		Restricted $\alpha_1 = \alpha_2$
	EUR/CHF	USD/CHF	EUR/CHF	USD/CHF	EUR/CHF USD/CHF
$\hat{\alpha}$	-1.8** (0.78)	-3.2** (1.3)	-2.2*** (0.77)	-3.2*** (1.2)	-2.6*** (0.84)
$\hat{\lambda}$	0.015** (0.0072)	0.014** (0.0071)	0.015** (0.007)		0.013* (0.0069)
Wald $H_0 : \alpha_1 = \alpha_2$ p-value			0.67 0.41		
J-statistic	0.28	0.99	1.1		1.7
p-value	0.59	0.32	0.78		0.79
No. Moment Conditions	3	3	6		6
No. Parameters	2	2	3		2
Note: The exchange rate variables are sampled at the close of the day (5:00 p.m. NY time). Please see notes of Table 3 for a detailed table description.					

Table C.2. OLS and IV Estimates of the Exchange Rate Response

	Δs^{EURCHF}	Δs^{USDCHF}
<i>OLS Estimates</i>		
$\Delta i^{\text{3M-LIBOR}}$	-0.4 (0.30)	2.5*** (0.50)
<i>IH-IV Estimates</i>		
$\Delta i^{\text{3M-LIBOR}}$	-2.2*** (0.65)	-2.4** (0.96)
<i>IH-IV Estimates with End-of-Day Exchange Rates</i>		
$\Delta i^{\text{3M-LIBOR}}$	-1.9*** (0.71)	-2.7*** (1.0)
<i>IH-IV Estimates with LIBOR Futures</i>		
Δi^{M90}	-1.2* (0.7)	-3.0*** (0.97)
Note: The table shows the OLS estimates. The policy rate variable is the three-month CHF LIBOR. *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively. The numbers in parentheses are standard deviations.		

Using these instruments, we estimate the parameters using two-stage least squares.

C.3 Results with End-of-Day Exchange Rates

The end-of-day exchange rates used for a robustness check are sampled at the close of the day, i.e., at 5:00 p.m. New York time. The IV-IH results are shown in Table C.2, whereas IH-GMM is displayed in Table C.1.

C.4 Results with Swiss LIBOR Futures

The ICE LIBOR futures data used for a robustness check is sampled at the close of the day, i.e., at 6:00 p.m. London time. The IV-IH results are shown in Table C.2, whereas IH-GMM is displayed in Table C.3.

Table C.3. GMM Estimates with LIBOR Futures

	Separate Equations		Joint Estimation		Restricted $\alpha_1 = \alpha_2$
	EUR/CHF	USD/CHF	EUR/CHF	USD/CHF	EUR/CHF USD/CHF
$\hat{\alpha}$	-1.5* (0.88)	-3.7*** (1.3)	-1.7* (0.91)	-3.7*** (1.3)	-3.1** (1.3)
$\hat{\lambda}$	0.012** (0.006)	0.012* (0.006)	0.011* (0.006)		0.0082 (0.0059)
Wald $H_0 : \alpha_1 = \alpha_2$ p-value			2.4 0.12		
J-statistic	0.84	1.3	1.5		4.0
p-value	0.36	0.25	0.68		0.41
No. Moment Conditions	3	3	6		6
No. Parameters	2	2	3		3
No. Iterations	3	3	3		3
Note: The policy rate variable is the change in the 90-day constant maturity rate computed from the CHF LIBOR futures. The exchange rate variables are sampled at the close of the day for futures (6:00 p.m. London time). Please see notes of Table 3 for a detailed table description.					

Table C.4. Non-policy Dates

19-Jan-2000	25-Jul-2002	14-Sep-2005	05-Nov-2008
02-Feb-2000	18-Sep-2002	14-Dec-2005	19-Nov-2008
22-Mar-2000	12-Dec-2002	15-Mar-2006	10-Dec-2008
14-Jun-2000	05-Mar-2003	14-Jun-2006	11-Mar-2009
13-Sep-2000	19-Mar-2003	13-Sep-2006	17-Jun-2009
07-Dec-2000	12-Jun-2003	13-Dec-2006	16-Sep-2009
21-Mar-2001	17-Sep-2003	14-Mar-2007	09-Dec-2009
13-Jun-2001	11-Dec-2003	13-Jun-2007	10-Mar-2010
10-Sep-2001	17-Mar-2004	12-Sep-2007	16-Jun-2010
21-Sep-2001	16-Jun-2004	12-Dec-2007	15-Sep-2010
06-Dec-2001	15-Sep-2004	12-Mar-2008	15-Dec-2010
20-Mar-2002	15-Dec-2004	18-Jun-2008	16-Mar-2011
01-May-2002	16-Mar-2005	17-Sep-2008	15-Jun-2011
13-Jun-2002	15-Jun-2005	07-Oct-2008	02-Aug-2011
Note: The table lists the non-policy sample dates. We follow Rigobon and Sack (2004) and use the previous day of the policy decision as non-policy day.			

C.5 Selection of Non-policy Sample

In our benchmark model we followed Rigobon and Sack (2004) and use the previous day of the policy decision as non-policy day (see Table C.4). The previous day is a valid candidate for the non-policy sample since no information has been released, market participants are in wait-and-see mode, and SNB officials are in the blackout period, when speaking publicly is not allowed. If we move more days away from the policy date, we get in conflict with monetary policy decision meetings by the Fed, which are typically more than two days prior to an SNB policy meeting. The assumption of homoskedastic shocks is potentially violated in these cases. In a further robustness exercise, we randomly select days from the candidate set of non-policy days and estimate our benchmark model each time. We find that the results remain qualitatively robust (i.e., a policy rate increase leads to an appreciation of the currency). Quantitatively, however, we observe that the median of the simulation results are absolute smaller than our benchmark (around -1.6 percent) and estimates are less stable. This finding is in line with the robustness results reported in Rigobon and Sack (2004). It could be the case

that the homoskedastic assumption, which is one of the key assumptions of the identification-through-heteroscedasticity approach, is violated because different kinds of shocks may materialize on these days.

References

- Adrian, T., R. K. Crump, and E. Moench. 2013. "Pricing the Term Structure with Linear Regressions." *Journal of Financial Economics* 110 (1): 110–38.
- Canetg, F., and D. Kaufmann. 2019. "Shocking Interest Rate Floors." Mimeo.
- Cuche-Curti, N., H. Dellas, and J.-M. Natal. 2009. "DSGE-CH: A Dynamic Stochastic General Equilibrium Model for Switzerland." SNB Economic Study No. 5.
- Ferrari, M., J. Kearns, and A. Schrimpf. 2017. "Monetary Policy's Rising FX Impact in the Era of Ultra-low Rates." BIS Working Paper No. 626.
- Fink, F., L. Frei, and O. Gloede. 2022. "Global Risk Sentiment and the Swiss Franc: A Time-Varying Daily Factor Decomposition Model." *Journal of International Money and Finance* 122 (April): Article 102539.
- Grisse, C. 2020. "The Effect of Monetary Policy on the Swiss Franc: An SVAR Approach." SNB Working Paper No. 2020-02.
- Grisse, C., and S. Schumacher. 2018. "Term Structure Dynamics at Low and Negative Interest Rates — Evidence from Switzerland." *Swiss Journal of Economics and Statistics* 154 (1): 1–17.
- Jarocinski, M., and P. Karadi. 2020. "Deconstructing Monetary Policy Surprises — The Role of Information Shocks." *American Economic Journal: Macroeconomics* 12 (2): 1–43.
- Kearns, J., and P. Manners. 2006. "The Impact of Monetary Policy on the Exchange Rate: A Study Using Intraday Data." *International Journal of Central Banking* 2 (4): 157–83.
- Kerssenfischer, M. 2019. "Information Effects of Euro Area Monetary Policy: New Evidence from High-frequency Futures Data." Discussion Paper No. 07/2019, Deutsche Bundesbank.
- Kugler, P. 2020. "The Short-run Impact of Interest Rates on Exchange Rates: Results for the Swiss Franc against the Euro

- and US Dollar from Daily Data 2001–2011.” WWZ Working Paper No. 2020/01.
- Lenz, C., and M. Savioz. 2009. “Monetary Determinants of the Swiss Franc.” SNB Working Paper No. 2009-16.
- Nitschka, T., and N. Mirkov. 2016. “Semi-parametric Estimates of Taylor Rules for a Small, Open Economy — Evidence from Switzerland.” *German Economic Review* 17 (4): 478–90.
- Ranaldo, A., and E. Rossi. 2010. “The Reaction of Asset Markets to Swiss National Bank Communication.” *Journal of International Money and Finance* 29 (3): 486–503.
- Rigobon, R. 2003. “Identification through Heteroskedasticity.” *Review of Economics and Statistics* 85 (4): 777–92.
- Rigobon, R., and B. Sack. 2004. “The Impact of Monetary Policy on Asset Prices.” *Journal of Monetary Economics* 51 (8): 1553–75.
- Rosa, C. 2011. “The High-frequency Response of Exchange Rates to Monetary Policy Actions and Statements.” *Journal of Banking and Finance* 35 (2): 478–89.
- Rudolf, B., and M. Zurlinden. 2014. “A Compact Open Economy DSGE Model for Switzerland.” SNB Economic Study No. 8.
- Soederlind, P. 2010. “Reaction of Swiss Term Premia to Monetary Policy Surprises.” *Swiss Journal of Economics and Statistics* 146 (1): 386–404.
- Zivot, E., and J. Wang. 2007. *Modeling Financial Time Series with S-PLUS*, Vol. 191. Springer Science & Business Media.

Demand or Supply? Price Adjustment Heterogeneity during the COVID-19 Pandemic*

Almut Balleer,^a Sebastian Link,^b Manuel Menkhoff,^c
and Peter Zorn^d

^aRWI Essen, TU Dortmund, IIES Stockholm, CESifo, and CEPR

^bifo Institute, LMU Munich, IZA, and CESifo

^cifo Institute and LMU Munich

^dLMU Munich and CESifo

We study price-setting behavior and subjective perceptions in German firm-level survey data to infer the relative importance of supply and demand during the COVID-19 pandemic. Demand shortages dominate at the onset of the pandemic. A reported negative impact of COVID-19 on current business is associated with a rise in the probability to decrease prices up to 10 percentage points in this episode. Supply forces gain in importance during the ensuing sudden surge in inflation and firms perceive goods supply shortages as most restrictive. Firms adversely affected during the early inflation decline show no higher probability of price increases.

JEL Codes: E31, E32, E60, D22.

*Previous versions of this paper were titled “Demand or Supply? Price Adjustment during the Covid-19 Pandemic” and appeared as preprints in *Covid Economics: Vetted and Real-Time Papers*, 31: 59–102. First Draft: June 18, 2020. We thank Klaus Adam (editor), anonymous referees, audiences at the Franco-German Fiscal Policy Seminar, European Central Bank, Norwegian School of Economics, 2021 ifo Conference on Macroeconomics and Survey Data, Birmingham Workshop on the Economic Response to Covid-19, 2021 European Economic Association, 2021 Verein für Socialpolitik, 2022 ICEA Inflation Conference for helpful comments, Martin Popp for sharing sectoral data on physical proximity to others at work, Jean-Victor Alipour for sharing sectoral data on work-from-home capacity, Martin Haas for excellent research assistance, the ifo Institute survey team for adding questions related to COVID-19 to the ifo Business Climate Survey, and the LMU-ifo Economics & Business Data Center (EBDC) for assistance with the data. Link and Menkhoff gratefully acknowledge funding from the German Bundesbank under the research grant “Long-COVID? Medium-Term Impact of the Corona Virus Pandemic on Decisions of Firms.” Declarations of interest: none. Corresponding author (Zorn): peter.zorn@econ.lmu.de.

1. Introduction

The 2020 COVID-19 recession and its aftermath has been disrupting economies across the globe. This paper estimates the relative importance of supply and demand in this episode, a key determinant for the scope and design of effective stabilization policy. To this end, we study producers' price-setting decisions along the extensive margin as well as their subjective perceptions about the nature of adverse effects due to the pandemic using German firm-level panel data from the ifo Business Climate Survey (ifo-BCS).

The results suggest that demand plays a dominant role to understand the early decline in economic activity in the wake of the COVID-19 pandemic. Relative to those weakly affected, firms that report a negative impact of COVID-19 on their current business situation are up to 10 percentage points—or three standard deviations—more likely to decrease prices. In a standard demand–supply framework, increasing prices reflect supply shortages and decreasing prices reflect demand shortages, all else equal. Similarly, firms attribute the most adverse effects of the pandemic to domestic and foreign demand through the spring of 2021. After that, supply forces gain in importance and firms perceive goods supply shortages as the most adverse effect. Firms that lowered their prices at the onset of the pandemic due to a more adverse impact do not show a relatively higher chance for price increases. Firms' subjective perceptions suggest persistent adverse effects of the pandemic at the firm level.

Consistent with these findings, aggregate producer price inflation in 2020 declined in the wake of the COVID-19 pandemic, followed by a sudden surge in 2021. However, firm-level panel data offer at least four advantages over aggregate time series: we can (i) control for other determinants of price adjustments that affect inflation irrespective of the COVID-19 outbreak, including the preceding decline in economic activity, (ii) overcome aggregation issues that potentially bias inflation dynamics,¹ (iii) track rapid changes in economic activity, given the monthly frequency of the data, and

¹Cavallo (2020) and Alvarez and Lein (2020) show that short-run shifts in expenditure patterns bias price indices of *aggregate* inflation. This challenge does not apply to individual price quotes.

(iv) link observed price-setting behavior to additional firm-level outcomes and firms' perceptions. For example, the survey questionnaire regularly asks firms about the role of (intermediate) goods supply and labor supply shortages. Another major advantage of firm-level panel data over time-series or sectoral data is that it enables us to investigate the occurrence of price adjustments across and within industries.

The empirical analysis uses the German ifo Business Climate Survey (ifo-BCS). The ifo-BCS is a monthly, mostly qualitative, firm-level survey among a representative sample of roughly 6,000 firms in key sectors of the German economy. Central for this paper are its regular questions about price setting and additional information on firm-specific economic activity that allow to control for several determinants of price-setting behavior. Specifically, we observe realized and three-months-ahead planned price changes along the extensive margin together with firms' current business situation and business expectations. Planned price changes are also available for goods and services that are temporarily not in trade due to COVID-19 containment measures. Our baseline analysis therefore studies planned price changes. We obtain similar results when we use realized price changes. Planned and actual price-setting behavior in the ifo-BCS relatively closely co-move with aggregate producer price inflation and the correlation with quantitative changes in industry-specific producer price indices is relatively high, considering the data only provide qualitative survey information on the extensive margin of price adjustment.² A supplement to the ifo-BCS contains questions related to COVID-19. Among other things, firms in the survey assess the impact of the pandemic on their current business situation and, infrequently, report their perceptions on the relative importance of supply and demand forces.

We show that price decreases underlying the decline in inflation occur broadly across industries. Because the frequency of price decreases in any given industry is low, the fact that a larger share of firms now decreases prices results in higher firm-level heterogeneity within industries. Arguably, this fact is more consistent with low demand (e.g., due to low sentiment or high uncertainty), for supply

²The micro data underlying the German producer price index are not available for this time period.

shortages in only some sectors would result in more heterogeneous effects across industries. Price increases during the sudden surge in inflation are relatively more concentrated in some industries.

We estimate the effects of COVID-19 on the probability to adjust prices in the ifo-BCS during the decline in producer price inflation in April and May 2020. Prior to the pandemic, firms that are differentially affected by COVID-19 display very similar dynamics in planned price changes. Our baseline specification includes sector-by-time fixed effects which flexibly control for industry-specific trends and heterogeneous effects across sectors—for example, due to sector-specific government regulations related to the pandemic or differential reliance on oil in conjunction with the sharp decline in oil prices. Despite this flexible approach to control for sector-specific differences in price-setting behavior during the pandemic, we find a substantial rise of up to 10 percentage points in the probability of planned price decreases for firms with a strongly negative impact relative to firms with no or only weak impact of COVID-19. This effect is economically large and implies substantial firm-level heterogeneity in price setting due to differential exposure to the pandemic within sectors. We also find a concurrent decline in the probability of planned price increases for these firms. Conversely, positively affected firms display an approximately 8 percentage point higher chance of planned price increases and are less likely to plan price decreases. Since more than half of the firms report negative effects due to COVID-19 while only 5 percent report positive effects, these findings suggest a key role for demand during the early decline of inflation.

To some extent, the price decreases in 2020 also reflect quality deterioration in the supply of certain goods and services due to added health risks (see, e.g., Eichenbaum, Rebelo, and Trabandt 2021). This is important for policy, since it is critical to differentiate between health risks and economic forces behind price decreases. We follow Mongey, Pilossoph, and Weinberg (2021) in accounting for firm exposure to health risks and use a contact intensity index based on O*NET survey data and a work-from-home capacity index by Alipour, Falck, and Schüller (2021).³ We find that health risks across sectors are associated with a higher chance of a price decrease,

³Kaplan, Moll, and Violante (2020) employ a similar distinction between regular and social sectors in the economy.

while the probability to decrease prices remains significantly elevated in 2020, consistent with additional demand forces weighing on price adjustment.

We then study the sudden surge in inflation. In 2021, the fraction of price increases strongly rises across firms to levels well above those before the pandemic. Firms that were adversely affected at the onset of the crisis do not drive this result. The probability to raise prices and reverse earlier price decreases is not significantly higher at the extensive margin in these firms. Consistent with this finding, firms' subjective perceptions reported in the survey suggest that supply shortages at the end of the sample are not associated with demand deficiencies at the onset of the pandemic. Hence, demand shortages and supply disruptions during the early decline in inflation appear to display persistent effects on the level of firm prices. (Intermediate) Goods supply shortages gain in relative importance over time, broadly across firms and in some sectors in particular, which drives the observed increase in between-industry variation. Labor supply shortages, oil exposure, and energy prices are not significantly correlated with price increases across sectors once we control for (intermediate) goods supply shortages, which strongly correlate with the frequency of price increases.

A growing empirical literature studies the impact of COVID-19 on firms, in particular on prices and demand.⁴ Behavioral responses in consumer expenditures possibly mirror heterogeneity in pricing to some extent. Jaravel and O'Connell (2020) find a spike in U.K. consumer price inflation at the onset of the COVID-19 pandemic. Interestingly, U.K. producer prices declined in this period, in line with our results. Similarly, Meyer, Prescott, and Sheng (2022) find that U.S. firms expect to decrease their output prices at the onset of the pandemic.

Balduzzi et al. (2020) show that credit constraints and deaths due to COVID-19 correlate with firms' planned price increases. Alekseev

⁴For example, see Bartik et al. (2020), Buchheim et al. (2022), and Hassan et al. (forthcoming) on the firm-level impact of COVID-19 and Baker et al. (2020), Cabral and Xu (2021), Carvalho et al. (2021), Cavallo (2020), and Chetty et al. (forthcoming) on prices and consumer spending.

et al. (2023) find evidence that deteriorating financial conditions lead to more price decreases. In a robustness exercise, we control for credit constraints and find a lower probability of price increases in credit-constrained firms. The main estimates for the effect of COVID-19 on price adjustment remain unchanged in this exercise.

A few studies directly address the relative importance of demand and supply in the pandemic. Brinca, Duarte, and Castro (2021) find that labor supply dominates labor demand at the onset of the pandemic. In line with our findings, Meyer, Prescott, and Sheng (2022) find that U.S. firms mainly assess the pandemic as demand-driven.

Existing theoretical and quantitative work on the pandemic concentrates on sectoral heterogeneity (e.g., Guerrieri et al. 2022). Our results suggest a complementary and important role for firm-level heterogeneity. We therefore link this literature to empirical work on firms in the pandemic by capturing the relative importance of firm-level and sectoral heterogeneity. Firm heterogeneity mostly matters for the decline in inflation and sectoral heterogeneity relatively more during its surge.⁵ Research on the transmission of the COVID-19 pandemic further highlights the importance of weak demand (e.g., Baqaee and Farhi 2022; Caballero and Simsek 2021; Eichenbaum, Rebelo, and Trabandt 2021; Guerrieri et al. 2022). Our finding that prices tend to decrease rather than increase at the beginning of the pandemic provides direct reduced-form empirical evidence consistent with this line of work. Lastly, the evidence for persistence in the firm-level impact of COVID-19 adds to similar findings for employment and revenues in U.S. firms (Barrero et al. 2021).

The remainder of this paper is organized as follows. Section 2 explains the data and descriptive statistics, including a variance decomposition of between- and within-industry heterogeneity. Section 3 studies the differential impact of the pandemic on firms and the role of health risks in the early decline in inflation. Section 4 presents results for the sudden surge in inflation. Section 5 concludes.

⁵These results are consistent with models of price adjustment in which idiosyncratic shocks dominate firm-level price-setting behavior (e.g., Nakamura and Steinsson 2010).

2. Data and Descriptive Statistics

The main data source is the ifo Business Climate Survey (ifo-BCS), a long-standing monthly survey among a large panel of German firms.⁶ We limit the analysis to the manufacturing, services, and retail/wholesale industries that cover approximately 5,500 firms per survey wave on average.⁷ Table A.1 in Appendix A compares the distribution of firms across size classes and industries in the survey to administrative data. The ifo-BCS provides a very good representation of the German economy along these dimensions.

The survey is mostly qualitative, including questions about firms' business situation and expectations, planned and realized price changes, as well as questions related to the supply and demand of goods and services. Since March 2020, the survey questionnaire includes supplemental questions related to the COVID-19 pandemic.⁸

The questions on price setting ask whether firms plan to increase, decrease, or leave unchanged their prices over the following three months, as well as a similar question on price realizations in the preceding month.⁹ The frequency of realized price changes in the

⁶See Sauer and Wohlrabe (2020) for extensive documentation of the ifo-BCS. The survey questionnaires are predominantly filled out by senior management such as firm owners, members of the executive board, or department heads (Sauer and Wohlrabe 2019). The ifo Business Climate Index, a widely recognized leading indicator of the German business cycle, is based on the ifo-BCS. More generally, Lehmann (2023) demonstrates high predictive power of aggregated ifo-BCS data for gross domestic product, industrial production, employment, investment, exports, and inflation.

⁷We leave out firms in construction and insurance due to a lack of comparability to the survey questionnaires in manufacturing, services, and retail/wholesale. Data harmonization across sectors follows Link (2020). In April and May 2020, the sample comprises on average 5,485 firms per month (1,941 in manufacturing, 1,937 in services, and 1,607 in retail/wholesale). Table A.2 in Appendix A provides more detailed information on the industry composition.

⁸Appendix C presents translations of all survey questions used in this paper.

⁹The ifo-BCS questions on planned and realized price changes are used in several articles. Bachmann et al. (2019) study the relation between uncertainty and price setting, Balleer, Hristov, and Menno (2020) investigate the link between financial constraints and price setting, Link (2019) examines the effect of the 2015 minimum wage introduction on firms' price setting, and Enders, Hünnekes, and Müller (2019) study the effect of monetary policy announcements on firms' planned price changes.


ifo-BCS is on average essentially the same as in administrative micro data underlying the German producer price index (Balleer and Zorn 2019). Planned and realized price changes in the ifo-BCS are informative about aggregate inflation. Figure A.1 in Appendix A documents a relatively high correlation over time of these measures with aggregate inflation rates reported by the Federal Statistical Office, considering the data provide only qualitative survey information on the extensive margin of price adjustment and conceptional differences between aggregate inflation and planned price changes.¹⁰ In general, the correlation is slightly higher for planned than for realized price changes and highest in manufacturing, with a correlation coefficient of 0.84. Figure A.2 in Appendix A documents the corresponding correlation at the level of two- and three-digit industries. Several industries exhibit a correlation above 0.6, especially when products are homogeneous within industry.¹¹

In the wake of the COVID-19 pandemic, certain goods and services were temporarily not available or transferable due to government regulations and supply-chain disruptions. As a result, their market prices were either not observed or might not yet have responded to the pandemic shock. The baseline analysis therefore considers planned price changes that are available for goods and services not in trade and that possibly include intended price responses to the pandemic. However, we obtain similar results when we use realized price changes or restrict our sample to businesses not affected by closure.

Figure 1 shows the frequency of planned and realized price increases from 2018:M1 to 2021:M8. Price increases display seasonal

¹⁰We use the producer price index (PPI) for manufacturing firms, the wholesale price index (WPI) for wholesale, and the retail price index (RPI) for retailers. The German Federal Statistical Office does not provide a monthly producer price index for services. The implications of our results for CPI inflation are unclear. In general, the correspondence between the NACE industry classification system used in the ifo-BCS and the COICOP classification used in the CPI is ambiguous (Addessi, Pulina, and Sallusti 2017; Ganglmair, Kann, and Tsanko 2021). Moreover, the number of firms in the ifo-BCS in industries with direct correspondence to the CPI is small.

¹¹Industries with a weak correlation often reflect a high diversity of product, e.g., the computer, electronic, and optical products industry.

Figure 1. Frequency of Planned and Realized Price Changes over Time

Note: The panels depict the frequency of planned (left) and realized (right) price changes in the ifo-BCS. The sample covers the manufacturing, retail/wholesale, and services industries. The sample period starts in 2018:M1 and runs through 2021:M8 for planned price changes and through 2021:M7 for realized price changes reported in the following month. The vertical line refers to March 2020, i.e., the month when the COVID-19 pandemic reached Germany.

patterns and rise at the turn of the year.¹² Reflecting an earlier-starting decline in economic activity, the frequency of price increases falls throughout 2019, while the frequency of price decreases rises gradually. Planned price decreases then spike in April 2020, after the COVID-19 pandemic hit Germany in March 2020. Similarly, realized price decreases also climb and peak shortly after. The timing of abrupt movements is in line with Buchheim, Krolage, and Link (2022), who show that the COVID-19 pandemic hit after most firms in the ifo-BCS filled in the March 2020 survey questionnaires.¹³ On June 6, 2020, the German government unexpectedly announced a temporary reduction of the value-added tax (VAT) rate effective July through December 2020. Although the survey questionnaire asks

¹²Nakamura and Steinsson (2008) document similar seasonal patterns in U.S. producer price micro data.

¹³In early March, only a few German counties were strongly affected by COVID-19. Subsequently, infection rates increased exponentially resulting in nationwide school closures on March 13 and a nationwide curfew on March 22. Buchheim, Krolage, and Link (2022) document that firms' business outlook decreased strongest after March 13. Since roughly three out of four survey respondents filled in their survey questionnaire before, April 2020 is the first month in which the majority of survey respondents report reactions to COVID-19.

about prices excluding the VAT, the corresponding survey guideline is hidden in a footnote. It is possible that some firms falsely report price decreases due to the VAT rate cut and price increases when it expires, which we cannot rule out. At the beginning of 2021, price increases skyrocket, way above the typical seasonal pattern, and continue to rise over the course of 2021.¹⁴

The initial downward price adjustments occur broadly across industries while the ensuing upward price adjustments are relatively more concentrated in a few large industries. We rely on the following decomposition to establish this result:


$$\text{var}(p_{ij}) \equiv \bar{p}(1 - \bar{p}) = \underbrace{\sum_j \frac{N_j}{N} \bar{p}_j (1 - \bar{p}_j)}_{\text{within}} + \underbrace{\sum_j \frac{N_j}{N} (\bar{p}_j - \bar{p})^2}_{\text{between}}, \quad (1)$$

where p_{ij} is a binary indicator for price adjustment by firm i in industry j , N_j denotes the number of firms in sector j , N is the total number of firms, $\bar{p}_j$ is the industry mean of p_{ij} , and $\bar{p}$ its unconditional mean. The notation suppresses a time subscript for convenience. The second equality follows because p_{ij} is the realization of a Bernoulli random variable. The decomposition is an application of the law of total variance. We apply it separately to price increases and price decreases for each month in the sample.

Equation (1) is useful because it allows to relate time variation in the frequency of price adjustment (the cross-sectional mean of p_{ij}) to shifts in the cross-sectional distribution (the cross-sectional variance of p_{ij}) arising from variation within and between industries. The first term is a weighted average of dispersion in price adjustment within industries and captures the degree of firm-level heterogeneity within industries. The second term is a weighted variance of price adjustment frequencies across industries and captures variation between industries. The between-variance moves if the change in price adjustment frequencies is proportional or heterogeneous across sectors. Changes in larger sectors enter with larger weight.

¹⁴In addition to seasonality and the VAT cut expiration, there were other one-off factors that possibly affected producer prices in January 2021. First, the German minimum wage increased in January 2021 from 9.35 euros to 9.50 euros and further to 9.60 in July 2021. Second, a consumer tax on carbon emissions was introduced.

Figure 2. Between-Industry Variation in Price Adjustment



Note: The left panel plots the share of variance explained by between-industry variation in planned price increases and decreases over the sample period, based on two-digit NACE industries and at least five observations per sector. The vertical line depicts 2020:M3. The right panel compares the average frequency of planned price increases across two-digit industries in 2020:M8 and 2021:M8. The size of each bubble represents the relative number of firms in each sector. The black line is the 45-degree line.

The left panel of Figure 2 plots the share of variance explained by between-industry variation over the sample period, based on two-digit NACE industries and at least five observations per sector. Between-industry variation accounts for about 5 percent of the variation in price increases and for a little less than 5 percent in price decreases through 2020:M3. There are no marked changes around the onset of the pandemic. Within-industry and between-industry variation both contribute proportionately to the changes in the frequency of price adjustment observed in Figure 1, while within-industry differences explain the bulk of it. Between-industry variation becomes disproportionately important during the sudden surge in inflation. Its relative share in the frequency of price increases triples, up to 15 percent.

To illustrate this result, the right panel of Figure 2 compares the frequency of planned price increases across two-digit NACE industries in 2020:M8 and 2021:M8. The size of each bubble represents the relative number of firms in each sector. The figure shows an increase in between-industry variation between 2020:M8 and 2021:M8. While the frequency of price increases is higher in almost all industries in

August 2021 compared to August 2020, the increase in the frequency of price increases is not uniform across industries and becomes more dispersed. For some industries, the frequency of price increases remains the same (those close to or on the 45-degree line). For other, relatively large industries the frequency of price increases rises particularly strongly: retail, wholesale, paper products, metal products, and electrical equipment industries.

These findings do not change if we restrict the sample to at least 20 observations per two-digit industry or use four-digit industries, in which case the relative shares are about 10 percent each prior to the pandemic and jump to around 25 percent during the sudden surge in inflation in the case of price increases. The results remain essentially the same if we restrict the sample to a balanced panel. We also observe similar patterns within manufacturing and retail, but not in services. This finding suggests that variation between services on the one hand, and manufacturing and retail on the other also accounts for rising between-variation. Lastly, these results also carry over to realized price changes.

Taken together, the initial downward price adjustments underlying the early decline in inflation occur broadly across industries. This finding is arguably more consistent with relatively high importance of low aggregate demand (e.g., due to low sentiment or high uncertainty), for negative supply shocks would result in more heterogeneous effects across industries. Section 3 further explores the determinants of rising firm-level heterogeneity as well as the role of industry-specific health risks in this episode. During the sudden surge, price increases become relatively more concentrated in some, relatively large sectors. This finding suggests an increasing relative importance of supply shortages. Section 4 further investigates the role of skilled labor shortages, energy costs, (intermediate) goods supply shortages, and oil exposure during this period and finds goods supply shortages across sectors to be important. Overall, however, the share of price adjustments explained by variation across industries remains below a third.

3. The Early Decline in Inflation

We focus on April and May 2020 to study price adjustment behavior underlying the early decline in inflation for three reasons: First, at

the beginning of the pandemic we can unambiguously tell whether a price is higher or lower relative to the pre-pandemic period. That is no longer possible as the pandemic evolves, as we do not observe the intensive pricing margin. Second, the decline in inflation is most pronounced in these months. Third, falsely reported price decreases due to the VAT rate cut announced in June possibly raise the frequency of price decreases across all firms and blur the analysis. However, in light of Figure 1, we expect this effect to be small, if anything, and uniform across firms.

We first infer the relative importance of demand and supply from firms' price adjustments behavior, relying on variation within industries. Thus, the analysis relates the observed rising firm-level heterogeneity during the early decline in inflation to the differential impact of the pandemic across firms. We then address health risks as a prominent narrative behind sector-level heterogeneity at the onset of the pandemic.

3.1 Price Adjustment Behavior

Since April 2020, the ifo-BCS asks firms to assess the impact of the COVID-19 pandemic on their business situation on a seven-point scale ranging from -3 ("negative") to $+3$ ("positive"). We refer to these *COVID-19 impact* categories as "strongly negative" (-3), "negative" (-2), "weakly negative" (-1), and "no impact" (0), with analogous labels for the positive categories.

Section B.1 in Appendix B provides extensive descriptive statistics on *COVID-19 impact*. We summarize the main findings. First, there is substantial firm-level heterogeneity in *COVID-19 impact* mirroring the large within-industry variation in price adjustments of Section 2. Second, more adverse *COVID-19 impact* is associated with worse business conditions, pessimistic business expectations, lower capacity utilization, and stronger expected revenue losses. We also find large heterogeneity in firms' price-setting behavior during the early decline in inflation that correlates with *COVID-19 impact*. As documented in Section B.2 in Appendix B, positively affected firms tend to increase their prices, while negatively affected firms tend to decrease their prices in all sectors, in particular in retail/wholesale. In industries in which the majority of firms

report low orders, negatively affected firms decrease prices more and increase prices less.

We now formally explore differences in price adjustments across *COVID-19 impact* categories during the early decline in inflation in April and May 2020 while controlling for other determinants of price-setting behavior. Our baseline analysis focuses on planned price changes. We obtain similar results when we use realized price changes or when we restrict our sample to open businesses, i.e., those not affected by closure due to lockdowns.¹⁵

First, we estimate the following regression, separately for each month-year t between 2018:M1 and 2020:M5:

$$Y_{i,t} = \delta_{-3} \mathbb{1}(Covid_{i,04/20} = -3) + \delta_{-2} \mathbb{1}(Covid_{i,04/20} = -2) + \delta_{\{2;3\}} \mathbb{1}(Covid_{i,04/20} = 2 \vee 3) + X'_{i,t-3} \beta + \gamma_s + u_{i,t}. \quad (2)$$

Here, $Y_{i,t}$ refers to an indicator for planned price increases or decreases over the following three months for firm i . The indicator variables group firms according to their *COVID-19 impact* in April 2020 as being strongly negatively affected ($Covid_{i,04/20} = -3$), negatively affected ($Covid_{i,04/20} = -2$), or positively affected ($Covid_{i,04/20} = 2 \vee 3$), while firms with weak or no *COVID-19 impact* serve as the control group.¹⁶


In addition, we include two-digit NACE industry fixed effects (γ_s), and separate indicators for positive and negative responses to the questions about business situation, business expectations, and orders, each lagged by three months and collected in $X_{i,t-3}$, to control for past economic activity of firms.

Figure 3 shows the time series of the frequency of planned price increases and decreases for each *COVID-19 impact* category, net of controls. In every month, the difference between each line relative to the group of firms with weak or no *COVID-19 impact* corresponds to the estimated coefficient δ_i , with $i = -3, -2, \{2; 3\}$, from Equation (2). The frequency-weighted mean of all lines in a given month equals the month's sample mean.¹⁷

¹⁵See Figure A.3 and Table A.4 as well as Table A.5 in Appendix A.

¹⁶We group the positive *COVID-19 impact* categories because of their low number of observations.

¹⁷See Yagan (2015) for a similar approach in a different context.

Figure 3. Planned Price Adjustment Before and During the COVID-19 Pandemic

Note: This figure shows the time series of the frequency of planned price decreases (left) and price increases (right) for each grouped *COVID-19 impact* category as of 2020:M4, net of controls. Firms are grouped according to their *COVID-19 impact* in April 2020. In every month, the difference between each line relative to firms with weak or no impact corresponds to the estimated coefficient, δ_i , $i = -3, -2, \{2, 3\}$ from Equation (2). The frequency-weighted average of all lines in a given month equals the month's sample average. The vertical red line refers to March 2020, i.e., the month when the COVID-19 pandemic reached Germany. Sample: 2018:M1–2020:M5.

Planned price changes display similar patterns across *COVID-19 impact* categories before the pandemic. The left panel of Figure 3 illustrates that the frequency of planned price decreases displays essentially identical dynamics across impact categories prior to 2020:M3, indicated by the first vertical red line, when measures to prevent the spread of COVID-19 were installed (see footnote 13). Likewise, the right panel shows that the frequency of planned price increases displays similar dynamics across impact categories prior to 2020:M3. This suggests that these similar trends would have continued in the absence of COVID-19.

However, we observe that price-setting behavior of firms in different *COVID-19 impact* categories is highly heterogeneous after 2020:M3. The frequency of planned price decreases skyrockets for strongly negatively affected firms, rapidly rises for firms with negative impact, and remains at similar levels for positively affected firms. There is no comparable spike in the frequency of planned price increases that would suggest upward price pressure at the beginning of the pandemic. The frequency of planned price increases remains at

similar levels for firms with positive impact and falls for those with (strongly) negative impact. Overall, this suggests a strongly disinflationary effect of *COVID-19 impact* and large differences across firms within industries.

Next, we exploit the panel dimension of the ifo-BCS and the timing of events to account for level differences, seasonality, and business cycle movements observable in Figure 3, i.e., slight upward and downward trends in planned price decreases and increases, respectively, consistent with the cooling of the German economy during this period. We estimate the following panel regression on the sample 2018:M1 to 2020:M5:

$$Y_{i,t} = \delta_{-3} \mathbb{1}(Covid_{i,t} = -3) + \delta_{-2} \mathbb{1}(Covid_{i,t} = -2) + \delta_{\{2;3\}} \mathbb{1}(Covid_{i,t} = 2 \vee 3) + X'_{i,t-3} \beta + \alpha_i + \gamma_{t,s} + u_{i,t}. \quad (3)$$


Month-year fixed effects ($\gamma_{t,s}$) at the level of two-digit NACE industries flexibly control for industry-specific trends—for example, due to differential reliance on oil in conjunction with the sharp decline in oil prices. They also control for heterogeneous effects during the COVID-19 pandemic across sectors. As before, we therefore rely on within-industry variation in pricing decisions due to *COVID-19 impact*. Relative to Equation (2), we set the grouped *COVID-19 impact* categories to zero for all observations prior to 2020:M4. Firm fixed effects absorb time-invariant characteristics in price-setting behavior.

Figure 4 presents estimation results for planned price decreases (blue; filled) and planned price increases (red; hollow). Panel A contains estimates based on a sample pooling firms in all industries, while panel B presents regression results when estimating Equation (3) separately for firms in manufacturing, retail/wholesale, and services.¹⁸

On the one hand, firms reporting a negative impact of COVID-19 tend to lower prices. In the pooled sample of panel A, the probability of planned price decreases spikes by 10 percentage points for firms strongly negatively affected, relative to the base category

¹⁸Panel A of Figure 4 refers to columns 3 and 6 of Table A.3 in Appendix A. Columns 1 and 4 of that table contain, for completeness, estimation results when only the *COVID-19 impact* category indicators are included in the regression.

Figure 4. Effects of *COVID-19 Impact* on Planned Price Adjustment



Note: This figure reports estimates from linear regressions of indicators for planned price decreases/increases on indicators for *COVID-19 impact* categories in the pooled sample (left) and separately by sector (right), based on Equation (3). *COVID-19 impact* measures the impact of COVID-19 on the current business situation on a seven-point scale from -3 (“negative”) to $+3$ (“positive”) in the ifo-BCS, which we group and label “Strongly Negative” (-3), “Negative” (-2), and “Positive” ($+2$ and $+3$), and the base category “Weak/No Impact” ($-1, 0$, or $+1$). Control variables include separate indicators for positive and negative responses to the questions about business situation, business expectations, and orders, all lagged by three months. In addition, we control for firm fixed effects and month-year fixed effects at the levels of two-digit NACE industries. Appendix C provides translations of all corresponding survey questions. Ninety-five percent confidence bounds are based on standard errors clustered at the firm level. Sample: 2018:M1–2020:M5. Columns 3 and 6 of Table A.3 in Appendix A present the corresponding numerical estimates.

of weak or no *COVID-19 impact*, net of controls. For negatively affected firms, the probability of planned price decreases rises by about 5 percentage points. By contrast, positively affected firms experience a decline in the chance of planned price decreases by about 2 percentage points. All estimates are statistically significant and economically very large compared to the unconditional two-digit sectoral frequency of planned price decreases of 3.7 percent and within-sector standard deviation equal to 3.4 percentage points in the period 2018:M1–2019:M12.

On the other hand, negative *COVID-19 impact* is associated with less frequent price increases. Strongly negatively affected firms display an approximately 4 percentage point lower chance of planned

price increases. By contrast, firms that report a positive impact on their business situation show an 8 percentage point rise in the probability of planned price increases, relative to the unconditional two-digit sectoral frequency of 21.2 percent and within-sector standard deviation equal 9.2 percentage points in 2018:M1–2019:M12.

Columns 2 and 5 of Table A.3 in Appendix A display results with month-year fixed effects instead of time-by-industry fixed effects (level differences across industries are absorbed by firm effects). The coefficients are very similar compared to our main specification in Figure 4. This result suggests that large firm-level heterogeneity within industry plays an important role for price-setting behavior during the early decline in inflation. This reflects the large within-industry variation in *COVID-19 impact* described in Section 2.

Panel B of Figure 4 shows that (strongly) negative *COVID-19 impact* is associated with a higher probability of planned price decreases across all sectors. The point estimate for the effect of positive *COVID-19 impact* on the chance of planned price increases is similar across sectors but less precisely estimated if we do not pool observations. A statistically significant decline in the probability of planned price increases for firms reporting a strongly negative impact is only observable in the retail/wholesale sector, while the drop in the probability of planned price decreases for positively affected firms is significant only in manufacturing. The upshot is that the main pattern of a higher chance of planned price decreases in (strongly) negatively affected firms, which represent more than half of all firms, is robust and not specific to any of the broad sectors we consider.

While the large majority of firms generally does not plan to change their prices (compare panel C of Table B.1 in Appendix B), prices become more flexible across all *COVID-19 impact* categories. Column 9 of Table A.3 in Appendix A shows that in firms with strongly negative impact, the probability of planned price changes increases by about 6 percentage points, reflecting the increased likelihood of planned price decreases. The same is true for firms negatively affected by COVID-19, which increase the chance of planned price changes by about 3 percentage points. The probability of planned price changes for firms with positive impact rises by about 5 percentage points, reflecting an increase in the probability of planned price increases. Again, these estimates are economically sizable compared

to the unconditional two-digit sectoral frequency of planned price change of 24.9 percent and within-sector standard deviation equal to 8.9 percentage points in the period 2018:M1–2019:M12.

These results speak to the relative importance of supply and demand shortages at the beginning of the pandemic. The result that negatively affected firms have a higher probability to decrease prices suggests that demand shortages dominate. In a basic demand–supply framework, a reduction in the supply of goods and services leads to increasing prices, holding demand constant. Conversely, given production, demand shortages lead to decreasing prices.

These results remain robust in five alternative specifications presented in Table A.6 in Appendix A. First, we weight the regressions by firm size to account for the relative importance of a firm in aggregate price indices.¹⁹ Second, we control for credit constraints. Financial frictions interact with price-setting behavior (Balleer, Hristov, and Menno 2020; Gilchrist et al. 2017; Kim 2020) and possibly also determine *COVID-19 impact* (Alekseev et al. 2023; Balduzzi et al. 2020). To measure credit constraints, we rely on a quarterly (at the end of each quarter) question in the ifo-BCS that asks whether firms negotiated for loans with banks over the past three months and, if so, the terms offered. We define firms as credit constrained if their most recent response was “restrictive” (in 2020:M4 and 2020:M5 we use the 2020:M3 value, for instance) and add this indicator as a control. Financial constraints significantly reduce the frequency of price increases, similar in size to *COVID-19 impact*. The estimated effects of *COVID-19 impact* remain virtually unchanged in this case.

Third, we include industry-by-time fixed effects at the four-digit instead of the two-digit industry level, which allows for differential effects within broader sectors but also reduces the number of observations slightly due to some cells being occupied by a single observation. Fourth, we restrict the sample to complete price spells and add indicator variables to control for Taylor pricing, i.e., price changes that occur in fixed time intervals (e.g., every six months;

¹⁹We construct firm weights following the procedure used by the ifo Institute in calculating the ifo Business Climate Index (see Sauer and Wohlrabe 2020 for details): The ifo-BCS micro data contain weights that are based on the number of employees in manufacturing and based on revenues in retail, wholesale, and services. The weights are scaled such that in a given industry their sum equals that industry’s share in gross value-added.

see Bachmann et al. 2019 and Lein 2010). Fifth, respondents may also consider realized price changes in answering the survey question on *COVID-19 impact*. Consequently, *COVID-19 impact* would be lower if prices decreased for a given change in output. This concern is partly alleviated by using planned price changes instead of realized price changes. Here, we estimate Equation (3) on a subsample of firms which did not change prices in the current month. In all of these alternative specifications, the main results remain robust.

3.2 Health Risks and Planned Price Adjustments

In this section, we investigate the role of health risks for price adjustments during the early decline in inflation. Health risks associated with physical contacts are a key characteristic of the COVID-19 pandemic. Adverse *COVID-19 impact* is associated with larger health risks (see Section B.1 of Appendix B). For policy, it is critical to differentiate between health risks and economic forces. If policymakers seek to contain the pandemic at minimal economic cost, stimulus policy to stabilize demand shortages related to health risks would backfire, for instance.²⁰


Yet, health risks vary mostly across sectors, and Section 2 shows that between-industry variation does not play a dominant role in price adjustments. Indeed, variation in health risks across sectors is only weakly positively related to the frequency of price decreases and very weakly negatively related to the frequency of price increases, as Section B.1 of Appendix B documents. We therefore expect the role of health risks for price changes during the early decline in inflation to be limited.

To explore this more formally, we split the sample by health risk exposure at the sectoral level. We use O*NET survey data to measure contact intensity to co-workers and customers at work at the five-digit industry level.²¹ Possible answer categories for the relevant

²⁰Fetzer (2022), for example, shows that a subsidy for food and non-alcoholic drinks in restaurants in the United Kingdom had a causal effect on higher COVID-19 infection rates.

²¹We thank Martin Popp for constructing this data by weighting the O*NET survey information on physical proximity with occupational employment per sector using the Integrated Employment Biographies (IEB) of the Institute for Employment Research (IAB) in Germany. More details about the O*NET data are available here: <https://www.onetonline.org/find/descriptor/result/4.C.2.a.3>.

Figure 5. Effects of Health Risk Exposure on Planned Price Adjustment



Note: The figure shows the time series of the frequency of planned price decreases (left) and planned price increases (right) by health risk exposure, net of controls. “High health risk exposure” refers to the group of firms in industries with above-median O*NET contact intensity and below-median capacity to work from home by Alipour, Falck, and Schüller (2021). The remaining firms are classified as “low health risk exposure.” The vertical red line refers to March 2020, i.e., the month when the COVID-19 pandemic reached Germany. See notes to Figure 3 for details on the construction of these figures. Sample: 2018:M1–2020:M5.

question range from 0 (“I don’t work near other people (beyond 100 ft.)”) to 100 (“Very close (near touching)”). We combine this measure with the share of workers in each two-digit industry able to work from home (WfH) provided by Alipour, Falck, and Schüller (2021) and define “high health risk exposure” by above-median contact intensity and below-median WfH capacity.²² The remaining observations are grouped as “low health risk exposure.”

First, we re-estimate Equation (2) replacing *COVID-19 impact* with this indicator for high health risk exposure. Figure 5 shows the observed trends in the probability of planned price decreases and planned price increases of high health risk exposure sectors parallel those of low health risk exposure sectors. In panel B, the probability of planned price increases is in general higher for high health risk exposure sectors. During the COVID-19 pandemic, the probability

²²Results are similar if we only use above-median O*NET contact intensity to define high health risk exposure. Table A.2 in Appendix A provides summary statistics for these COVID-19 infection risk proxy variables by two-digit industry.

of planned price decreases rises and the probability of planned price increases drops in both groups, but relatively stronger so in high health risk exposure sectors.

Next, we quantify the differential effect of high risk exposure during the COVID-19 pandemic in a difference-in-differences regression:

$$\begin{aligned} Y_{i,t} = & \delta_1 \mathbb{1}(\text{HighHealthRisk}_i = 1) + \delta_2 \mathbb{1}(\text{CrisisPeriod} = 1) \\ & + \delta_3 \mathbb{1}(\text{HighHealthRisk}_i = 1) \times \mathbb{1}(\text{CrisisPeriod} = 1) \\ & + X'_{i,t-3} \beta + u_{i,t}. \end{aligned} \quad (4)$$

Again, $Y_{i,t}$ denotes an indicator for a planned price decrease or a planned price increase of firm i at time t , controlling for past economic activity ($X_{i,t-3}$). The coefficient δ_1 captures any level differences in the probability to change prices between high and low health risk exposure before the COVID-19 pandemic. The coefficient δ_2 measures a level effect on the probability to change prices of low health risk exposure firms during the COVID-19 pandemic after 2020:M3. The coefficient δ_3 describes the differential effect on high health risk exposure firms.

Table 1 presents results. Column 1 shows that high health risk exposure firms are about 5 percentage points more likely to decrease their price during the COVID-19 pandemic relative to low health risk exposure firms. The probability of a planned price decrease remains significantly elevated by about 6 percentage points in low health risk exposure firms, consistent with the notion of firm-level heterogeneity and additional demand forces affecting planned price-setting behavior during the COVID-19 pandemic. Column 3 shows that the probability of planned price increases drops during the COVID-19 pandemic and is not statistically different across low and high health risk exposure.²³

As alternatives to the health risk exposure measure constructed from the O*NET contact intensity index and the WfH capacity

²³We also investigated spillover effects between low and high health risk firms through the input-output network. To do so, we added an interaction term of health risks in two-digit industries upstream or downstream, weighted by their input or output shares, and a crisis period indicator, without any effect on the reported estimates.

Table 1. Health Risks Effects of the COVID-19 Pandemic on Planned Price Adjustment

	Planned Price Decrease		Planned Price Increase	
	(1)	(2)	(3)	(4)
Crisis	0.057*** (0.0073)	0.069*** (0.012)	-0.091*** (0.011)	-0.099*** (0.0100)
High Health Risk	-0.000067 (0.0079)		0.043* (0.025)	
Contact Intensive (Eichenbaum et al. 2020)		0.0094 (0.0093)		0.095*** (0.012)
High Health Risk $\times$ Crisis	0.050*** (0.015)	0.011 (0.027)	-0.014 (0.021)	0.011 (0.016)
Contact Intensive (Eichenbaum et al. 2020) $\times$ Crisis				
Observations	121,099	121,111	121,099	121,111
Industry FE	No	No	No	No
Controls	Yes	Yes	Yes	Yes

Note: This table reports estimates from difference-in-differences regressions based on Equation (4). Firms are classified into high and low health risk groups based on three different measures: columns 1 and 3 group firms to “high health risk exposure” if they operate in industries with above-median contact intensity as described in footnote 21 and below-median industry-specific capacity to work from home by Alipour, Falck, and Schüller (2021). In columns 2 and 4, firms are sorted according to the heuristic classification by Eichenbaum et al. (2020) to “high contact goods and services.” Control variables include separate indicators for positive and negative responses to the questions about business situation, business expectations, and orders, all lagged by three months. Appendix C provides translations of all corresponding survey questions. Standard errors in parentheses are clustered at the firm level. Sample: 2018:M1–2020:M5. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

index, we also consider the heuristic classification of “high contact goods and services” by Eichenbaum et al. (2020). Columns 2 and 4 report estimates using this alternative measure and show that the probability of a planned price decrease increases during the COVID-19 pandemic. However, the effect on high health risk exposure firms is not statistically different from those of low health risk exposure firms in the case of price decreases. These firms are significantly more likely to raise their prices during the early decline in inflation, though.


In sum, our results suggest that exposure to health risks plays a significant role for price adjustment between sectors. However, a substantial part of the decline in price adjustments cannot be associated with health risks.

4. The Sudden Surge in Inflation

We now study the sudden surge in inflation in 2021. First, are those firms that were adversely affected at the beginning of the pandemic and that decreased their prices more likely to raise their prices in 2021? Second, does the relative importance of supply and demand shift over time? To this end, we study firms’ subjective perceptions about the adverse effects of the pandemic given in the survey. Third, the variance decomposition in Section 2 shows that, unlike the early decline in inflation, this episode is associated with an increase in the relative importance of variation across industries up to a third. What explains this increase? We study industry variation in intermediate input supply shortages, labor supply shortages, and exposure to oil and energy prices.

4.1 Price Adjustment Behavior

Figure 6 shows how *COVID-19 impact* affects price-setting behavior through 2021:M8. Here, we re-estimate Equation (3) replacing current *COVID-19 impact* with its average between 2020:M4 and 2021:M3. We use the average of *COVID-19 impact* to fix the composition of firms and study whether these firms are more likely to raise prices during the sudden surge. Fixing the composition does

Figure 6. Planned Price Adjustment until August 2021

Note: This figure shows the time series of the frequency of planned price decreases (left) and price increases (right) for each grouped COVID-19 impact category net of controls. Firms are grouped according to their average value of *COVID-19 impact* between April 2020 and March 2021, rounded to the next integer. In every month, the difference between each line relative to firms with weak or no impact corresponds to the estimated coefficient δ_i , $i = -3, -2, \{2, 3\}$ from Equation (2). The frequency-weighted average of all lines in a given month equals the month's sample average. The vertical red lines refer to March 2020, i.e., the month when the COVID-19 pandemic reached Germany, and July 2020 and January 2021, when the value-added tax was temporarily decreased and reverted to the original level, respectively. Sample period: 2018:M1–2021:M8.

not affect the results, as *COVID-19 impact* is very persistent over time.²⁴

At the onset of the crisis, these effects compare to those we estimate in the previous section shown in Figure 4. Negatively affected firms plan to decrease their prices more often throughout the pandemic. The effect on the frequency of planned price decreases spikes at the onset of the pandemic, then declines and rises again towards the end of 2020 when the second wave of the pandemic hits Germany. The frequency to decrease prices drops towards the end of the sample.

The surge in price increases is not primarily driven by firms adversely affected at the beginning of the COVID-19 pandemic that

²⁴Figure A.4 in Appendix A shows that the share of firms within each *COVID-19 impact* category is persistent over time. Relative to Table B.1 in Appendix B, the average relationship between *COVID-19 impact* and business conditions, business expectations, capacity utilization, expected revenue changes, and price-setting behavior in the full sample is essentially unchanged.

increase their prices more often in 2021. Even though the estimated probability to decrease prices converges across *COVID-19 impact* groups, negatively affected firms still decrease prices more often than positively affected firms. At the beginning of the pandemic, the frequency of planned price increases falls across all *COVID-19 impact* categories bar positively affected firms, bottoms out in 2020:M7, and starts to strongly climb in 2020:M9 across firms to levels well above those before the pandemic. Still, positively affected firms continue to be more likely to revise prices upwards than weakly and negatively affected firms.²⁵ A similar pattern emerges when we directly condition on pricing decisions early in the pandemic independent from *COVID-19 impact*. Only a third of the firms that decreased their prices in April 2020 increase their prices in June, July, or August 2021. The corresponding frequency of those firms that increased their prices in April 2020 is twice as high.²⁶

One caveat of our analysis is that we do not observe the intensive margin of price adjustment. It is possible that adversely affected firms do not increase prices more often, but when they do, they increase prices more substantially. Yet, a substantial share of more than 60 percent of adversely affected firms chooses no price increases at all.

The VAT rate cut in 2020:M6 does not appear to impede the analysis. There are no abrupt changes in the differential effects across *COVID-19 impact* categories, the variation which the estimation strategy we pursue relies upon. To the extent that some firms falsely report price changes, these will be most likely absorbed by the time fixed effect.

4.2 *Firms' Perceptions*

In addition to differences in pricing across *COVID-19 impact*, Figure 6 shows that price increases become more frequent overall, speaking to the increasing relative importance of supply. We collect additional direct descriptive evidence for the extended sample which corroborates the relative importance of supply. In June 2020, November 2020, February 2021, and June 2021, the ifo-BCS added

²⁵Figure A.3 in Appendix A documents the estimated effect on realized prices.

²⁶See Figure A.5 in Appendix A.


a question asking managers about the subjective reasons for the adverse effects of COVID-19 (SQ6 in Appendix C). Specifically, the question assesses the adverse effects of financing conditions, labor input, supply-chain disruptions (goods supply), government containment regulations, and demand (domestic and foreign, separately) due to the COVID-19 crisis. Possible answer categories range from “no effect” (+1) to “large adverse effect” (+5). To capture their relative importance at the firm level at a given point in time, we compare each subjective reason relative to the firm mean of all those remaining. Figure 7 shows the relative importance of the three predominant subjective reasons for the pooled sample and separately for each sector. Positive values represent relatively important reasons and negative values relatively unimportant reasons. Overall, deficiencies in domestic and foreign demand dominate as the most important determinant for adverse economic effects through the spring of 2021. This independent evidence reinforces the relative importance of demand inferred from price adjustment behavior in the previous section. After that, supply shortages become the most adverse effect. Patterns are qualitatively similar across sectors. Demand shortages, especially in foreign demand, are perceived as more important among manufacturing firms, while goods supply shortages are more important in retail/wholesale and less in services.

Lastly, demand deficiencies at the onset are not related to supply shortages later on in the pandemic. Figure A.6 in Appendix A relates the relative importance of demand and supply, respectively, in June 2020 to June 2021 net of sectoral fixed effects and controls. While the relative importance of supply increased for all industries, there is no apparent shift from relative demand to relative supply within industry.

4.3 Sectoral Outcomes

Figure 8 relates planned price increases to potential sources of upward pressures on prices across sectors. The left column plots the share of firms with planned price increases in a given sector against the share of firms perceiving (skilled) labor and (intermediate) goods shortages. The right column plots the share of firms with planned price increases against the share of energy costs and oil exposure in total production in a given sector. All panels aggregate at the


Figure 7. Subjective Reasons for Adverse Economic Effects of COVID-19



Note: This figure shows the relative subjective importance of domestic and foreign demand as well as goods supply as reasons for the adverse economic effects of COVID-19 on firm business activity. Supplementary questions in the ifo-BCS ask firms to assess the adverse effects due to COVID-19 of financing conditions, demand (domestic and foreign, separately), labor input, supply-chain disruptions, and government containment regulations. Appendix C provides the corresponding translated survey questions. Possible answers categories range from “no effect” (+1) to “large adverse effect” (+5). To capture their relative importance at the firm level at a given point in time, we compare each subjective reason relative to the firm mean of all those remaining, and then average across firms. Thus, positive values reflect relatively more important reasons. Subjective reasons data are available in June 2020, November 2020, February 2021, and June 2021. Panel A covers the total sample; panels B through D cover the manufacturing, retail/wholesale, and services industries, respectively.

two-digit NACE industry level and use data on planned price increases from July 2021. We observe skilled labor supply shortages and oil price exposure for the entire economy, (intermediate) goods supply shortages in manufacturing and retail/wholesale, and energy costs only in manufacturing. The size of each bubble represents the

Figure 8. Planned Price Increases on Sector Level in July 2021



Note: This figure shows the share of planned price increases in relation to goods supply shortage, skilled labor supply shortage, oil price exposure, and energy costs at the two-digit NACE industry level in July 2021. Goods supply shortage is the share of firms in a given sector that reports to be constrained by lack of material/intermediates. Skilled labor supply shortage is the share of firms in a given sector that reports to be constrained by lack of skilled labor. Both measures derive from Q7 in Appendix C. Oil price exposure is the share of oil and oil-related (measured as the input share of cokery and oil products times the share of oil input in that sector) inputs in total production from the 2016 input-output matrix. We winsorize the oil exposure share in four sectors at 1. Energy costs is the total share of energy in gross production value. Both statistics are obtained from the Federal Statistical Office. Goods supply shortages are only available for manufacturing and retail/wholesale sectors. Energy costs are only available for the manufacturing sectors. The linear fit shown is from Table A.7 in Appendix A.

relative number of firms in each sector. We label the five sectors from Section 2 that contribute most to the rise in between-variance.

(Intermediate) Goods supply shortages are positively associated with price increases during the surge in inflation. Table A.7 in Appendix A contains univariate and joint regression results for the share of planned price increases on each of the variables shown. The

negative relation between price increases and labor shortages is statistically insignificant, and labor shortages in sectors that drive the surge in between-variance do not stand out. The share of energy costs in general does not significantly relate to price increases. The sectors that drive the sudden surge in inflation rather tend to have large oil exposure and in particular perceive (intermediate) goods supply shortages. In a multivariate regression on all variables, (intermediate) goods supply shortages emerge as the single significant explanatory variable for planned price increases. A 10 percentage point increase in the share of firms perceiving (intermediate) goods supply shortages is associated with a rise in the share of planned price increases by 4.6 percentage points. These results suggest that supply-chain disruptions in some sectors drive the increase in between-variance. Consistent with this evidence, we find that indicators for retail/wholesale and manufacturing alone can explain 67 percent of the sectoral differences (see column 8 in Table A.7). Hence, prices for goods and services diverge strongly. Rising sectoral heterogeneity in turn accounts for up to a third of the observed price adjustments during the sudden surge in inflation (see Section 2).

5. Conclusion

We infer the relative importance of demand and supply during the COVID-19 pandemic by studying price-setting behavior and subjective perceptions in German firm-level survey data. The estimates presented in Figure 4 imply that, at the onset of the pandemic, (strongly) negative *COVID-19 impact* is associated with a 5 (10) percentage point increase in the probability of planned price decreases, relative to weak or no *COVID-19 impact*. Differences in health risks across sectors are important, but do not account for the substantial heterogeneity in price adjustment behavior within industries. Through the lens of a simple demand–supply framework, these results thus suggest an important role for demand shortages, driven by economic forces, at the onset of the COVID-19 pandemic. Otherwise, a reduction in supply would have reversed the observed price-setting behavior, with higher probability of price increases associated with negative *COVID-19 impact*. The fact that survey respondents report demand shortages as the primary reason for the

adverse effects due to COVID-19, both at the onset of the pandemic and as it unfolds, provides additional, more direct evidence in support of this result. Our results provide reduced-form evidence for theoretical channels that highlight the role of demand deficiencies in the pandemic and suggest a role for policy to stabilize aggregate demand while containing the COVID-19 pandemic.


The fact that within-industry variation accounts for the bulk of price adjustments during the early decline in inflation is no contradiction to the body of theoretical and quantitative work that considers the role of sectoral heterogeneity, but highlights the relative importance of firm heterogeneity. The propagation channels stressed by this literature (e.g., complementarities in consumption or input-output networks) possibly also operate within industries. Additional reasons for within-sector heterogeneity—for instance, management skills or other business characteristics such as online representation or service (see Bloom, Fletcher, and Yeh 2021 for direct evidence)—need to be considered as well. This result is important for policy, as it advises caution against targeted industry-specific stimulus to buffer the COVID-19 shock.

Upward price adjustments during the sudden surge in inflation are not due to adversely affected firms that lowered their prices at the beginning of the pandemic. Instead, we find broad-based (intermediate) goods supply shortages across firms as well as some heterogeneity in industry exposure to account for the sudden surge in inflation. A key question for inflation dynamics therefore is the persistence of goods supply shortages. Labor supply shortages, energy prices, and oil exposure are not significantly correlated across sectors once we control for goods supply shortages.

Another potentially important determinant of price adjustment are firms' inflation expectations. Ongoing data collection efforts show that firms' inflation expectations in our sample have substantially increased from below 2 percent in 2020:M12 to over 3.5 percent in 2021:M9 on average (see Figure A.7 in Appendix A). They therefore potentially contribute significantly and substantially to the surge in price increases. At this moment, we are not able to link these expectations to the firm-level pricing decisions in our sample and investigate this channel more closely. We leave this exploration to future work once data linkage becomes possible.


Appendix A. Supplementary Figures and Tables

Figure A.1. Correlation of Planned and Realized Price Changes and Aggregate Inflation




Note: The figure plots time series of the monthly realized change in (seasonally adjusted) aggregate price indices (i.e., the producer price index (PPI) for manufacturing, the wholesale price index (WPI) for wholesale, and the retail price index (RPI) for retail) from the German Federal Statistical Office (Destatis); dashed blue line, right axis) relative to three months before against the frequency of price adjustment published by the ifo Institute (solid black line, left axis) based on planned (first row) and realized (second row) price changes from the ifo-BCS. Realized prices for month t are measured at the end of the preceding month and shown as the moving average of the two preceding months and the current month. The ifo series are seasonally adjusted and aggregated as a weighted mean using ifo firm size weights to aggregate up to NACE two-digit level industries and using gross value-added from industry to aggregate level. We show time series separately for manufacturing, wholesale, and retail (including car sellers) industries. Destatis does not provide a monthly producer price index for services, hence not displayed here. The vertical red line refers to 2020:M3.

Figure A.2. Correlation of Price Changes and Sectoral Inflation for Selected Industries




Note: The figure shows correlations between the monthly realized change in (seasonally adjusted) producer price indices from the German Federal Statistical Office (Destatis) relative to three months before and the planned (left column) and realized (right column) price changes from the ifo-BCS. Realized prices for month t are measured at the end of the preceding month and shown as the moving average of the two preceding months and the current month. The ifo series are seasonally adjusted and aggregated to the three-digit and two-digit industry level as a weighted mean using ifo firm size weights. Sectors with more than 30 observations per period only. Sample: 2007:M1–2021:M7.

Figure A.3. Realized Price Adjustment before and during the COVID-19 Pandemic




Note: This figure shows the time series of the frequency of realized price decreases (left) and price increases (right) for each grouped *COVID-19 impact* category net of controls. Firms are grouped according to their average *COVID-19 impact* between April 2021 and March 2020. In every month, the difference between each line relative to firms with weak or no impact corresponds to the estimated coefficient δ_i , $i = -3, -2, \{2, 3\}$ from Equation (2). The frequency-weighted average of all lines in a given month equals the month's sample average. The vertical red lines refer to March 2020, i.e., the month when the COVID-19 pandemic reached Germany, and July 2020 and January 2021 when the value-added tax was temporarily decreased and increased back to the original level, respectively. Sample: 2018:M1–2021:M7.

Figure A.4. COVID-19 Impact over Time




Note: This figure plots the shares of firms by *COVID-19 impact* over time. *COVID-19 impact* measures the impact of COVID-19 on the current business situation on a seven-point scale from -3 (“negative”) to $+3$ (“positive”) in the ifo-BCS, which we group and label “Strongly Negative Impact” (-3), “Negative Impact” (-2), “Weak/No Impact” ($-1, 0$, or $+1$), and “Positive Impact” ($+2$ and $+3$), respectively. Appendix C provides a translation of the corresponding survey question.

**Figure A.5. Planned Price Increases in 2021
Grouped by Price Plans in April 2020**




Note: This figure shows the frequency of planned price increases in June, July, and August 2021. Firms are grouped according to their planned price changes in April 2020.

**Figure A.6. Relative Importance of
Demand and Supply over Time**



Note: This figure compares the relative subjective importance of domestic demand and goods supply in June 2020 and June 2021 net of sectoral fixed effects and the same controls as in Equation (2). Supplementary questions in the ifo-BCS ask firms to assess the adverse effects due to COVID-19 of financing conditions, demand (domestic and foreign, separately), labor input, supply-chain disruptions, and government containment regulations. Appendix C provides the corresponding translated survey questions. Possible answers categories range from “no effect” (+1) to “large adverse effect” (+5). To capture their relative importance at the firm level at a given point in time, we compare each subjective reason relative to the firm mean of all those remaining, and then average across firms. Thus, positive values reflect relatively more important reasons.

Figure A.7. Inflation Expectations of Firms in the ifo-BCS over Time



Note: This figure plots (A) the sample mean and (B) the median of expected CPI inflation in Germany 12 months ahead. The inflation expectations are elicited on a quarterly basis in a new survey module among the firms in the ifo-BCS, which, however, cannot be linked to firms’ pricing decisions in the main survey that we use in this paper, yet.

Table A.1. Distribution of Firms Responding to the ifo-BCS by Industry and Size

Industry	ifo Business Climate Survey				Distribution of German Firms by		
	Small (1)	Medium (2)	Large (3)	Total (4)	Count (5)	Employees (6)	Value- Added (7)
Manufacturing	12.20	14.93	9.45	36.58	9.65	28.56	37.12
Energy, Water, and Waste	0.53	0.21	0.08	0.82	3.77	2.40	5.10
Retail, Wholesale, and Repair of Motor Vehicles	22.20	6.47	1.92	30.59	26.11	22.85	18.87
Transportation and Storage	1.94	1.37	0.72	4.02	4.81	8.21	6.57
Accommodation and Food Services	2.98	1.14	0.09	4.21	10.73	8.77	2.97
Information and Communication	3.38	1.63	0.47	5.48	5.78	5.27	7.69
Real Estate Activities	0.95	0.30	0.06	1.31	7.26	1.99	4.13
Professional, Scientific, and Technical Activities	9.09	2.24	0.44	11.76	22.09	9.61	9.75
Administrative and Support Services	2.94	1.39	0.89	5.22	9.80	12.34	7.80
Total	56.20	29.68	14.12	100.00	100.00	100.00	100.00
Distribution of German Firms by							
Count	96.48	2.79	0.73	100.00			
Employees	36.38	16.88	46.74	100.00			
Gross Value-Added	23.64	15.39	60.97	100.00			

Note: This table compares the distribution of firms in the ifo-BCS to the distribution of firms in administrative data by industry and firm size. The ifo-BCS data are based on the 2020:M4 survey wave. The administrative data are based on the 2019 Statistics on Small and Medium-sized Enterprises (*“Statistik für kleine und mittlere Unternehmen”*) provided by the Federal Statistical Office (EVAS Code 48121). The definitions of size classes are small: 0–49 employees; medium: 50–249 employees; large: 250+ employees.

Table A.2. Proxy Variables for Risk of COVID-19 Infection by Two-Digit Industry

Two-Digit NACE Industry	COVID-19 Impact	Contact Intensity Index	Work-from-Home Capacity	High Contact Intensity	Observations
10 Food Products	-0.66	58.92	41.83	0	148
11 Beverages	-2.42	57.31	49.41	0	55
12 Tobacco	NA	55.43	54.56	0	NA
13 Textiles	-1.77	53.84	57.91	0	70
14 Wearing Apparel	-2.25	56.18	65.42	0	28
15 Leather	-2.58	56.56	57.8	0	19
16 Wood	-1.28	56.25	42.82	0	149
17 Paper Products	-0.76	56.58	58.59	0	143
18 Printing	-1.86	52.32	63.42	0	131
19 Coke & Petroleum Prod.	-1.73	54.66	61.24	0	11
20 Chemicals	-1.05	54.18	60.79	0	209
21 Pharmaceuticals	0.22	54.1	62.59	0	23
22 Plastic	-1.31	56.55	48.28	0	274
23 Glass	-1.19	55.22	43.61	0	217
24 Basic Metals	-1.98	56.18	43.36	0	189
25 Metal Products	-1.75	56.89	43.33	0	639
26 Computer, Electronic, & Optical Prod.	-1.48	54.07	64.73	0	224
27 Electrical Eq.	-1.56	55.66	61.29	0	325
28 Machinery	-1.81	55.97	55.36	0	727
29 Cars	-2.35	56.32	55.12	0	125
30 Oth. Vehicles	-2.37	57.44	57.57	0	8
31 Furniture	-2.5	56.38	45.58	0	88
32 Oth. Manuf.	-1.73	54.59	54.29	0	62
33 Repair of Machinery	-2	57.26	55.64	0	10
35 Electricity	-1	57.06	68.43	0	7
37 Sewerage	-1	56.13	50.21	0	13
38 Waste	-0.79	54.62	42.31	0	67
39 Remediation	-0.67	56.98	48.08	0	6

(continued)

Table A.2. (Continued)

Two-Digit NACE Industry	<i>COVID-19 Impact</i>	Contact Intensity Index	Work-from-Home Capacity	High Contact Intensity	Observations
41 Constr. Serv.	NA	53.85	40.68	0	NA
43 Constr. Serv.	-1.23	61.86	42.39	0	22
45 Car Sellers	-2.12	57.67	52.1	0	197
46 Wholesale	-1.2	56.27	64.06	1	1,549
47 Retail	-1.33	63.79	46.84	1	1,468
49 Land Transp.	-1.69	59.86	29.85	0	170
50 Water Transp.	-2.5	62.29	49.46	1	16
51 Air Transp.	NA	68.2	43.04	1	NA
52 Warehousing	-1.79	57.72	42.73	0	216
53 Postal Serv.	-1.48	62.5	38.54	0	25
55 Hotels	-2.97	62.92	42.77	1	317
56 Restaurants	-2.87	69.77	38.98	1	118
58 Publishing	-1.76	53.32	83.62	0	74
59 Movie & TV Production	-2.19	58.68	69.84	0	48
60 Broadcasting	-2.52	53.68	79.77	0	21
61 Telecommunication	-0.3	53.2	73.49	0	23
62 Programming	-1.04	50.45	87.43	0	392
63 Information Serv.	-0.89	52.69	82.17	0	56
64 Financial Serv.	-1.87	51.47	89.43	0	69
65 Insurance	NA	53.01	89.55	0	NA
66 Oth. Fin.	-1.53	51.4	88.21	0	45
68 Real Estate	-1.1	50.57	70.93	0	142
69 Legal Serv.	-0.54	48.12	87.59	0	115
70 Consulting	-1.34	50.03	76.53	0	194
71 Architects	-0.96	50.23	70.69	0	700
72 Sc. Research	-1.08	51.46	70.56	0	110
73 Advertising	-1.91	52.19	78.08	0	108
74 Oth. Prof., Sc., & Techn. Serv.	-1.44	53.57	70.79	0	71

(continued)

Table A.2. (Continued)

Two-Digit NACE Industry	COVID-19 Impact	Contact Intensity Index	Work-from- Home Capacity	High Contact Intensity	Observations
77 Rental & Leasing	-1.85	58.05	59.98	0	55
78 Empl. Act.	-2.26	58.76	44.37	0	136
79 Travel Agencies	-2.84	58.05	55.64	0	116
80 Security & Inv. Act.	-0.81	71.59	42.63	0	32
81 Facility Man.	-1.12	51.97	38.15	0	69
82 Oth. Admin.	-2.12	57.28	71.86	0	136
84 Public Admin.	NA	55.74	69.26	0	NA
85 Education	-2.23	61.12	72.49	1	73
86 Health Serv.	-0.67	81.23	45.77	1	6
87 Resid. Care	NA	77.44	47.59	0	NA
88 Social Work	NA	78.22	52.48	0	NA
90 Arts	-2.4	65.8	63.51	1	20
91 Libraries	NA	60.14	65.94	0	NA
92 Gambling	-2.35	65.91	47.74	0	17
93 Sports	-2.69	66.67	61.61	1	32
94 Employers Assoc./Unions	NA	55.24	69.98	1	NA
95 Repair of Computers	-1.27	55.92	59.55	0	11
96 Oth. Personal Serv.	-2.37	62.52	41.76	0	8

Note: This table provides summary statistics of the data set used in this paper at the level of two-digit NACE industries. *COVID-19 impact* measures the sample average impact of COVID-19 on the current business situation which is elicited on a seven-point scale from -3 (“negative”) to +3 (“positive”) in the ifo-BCS. The “Contact Intensity Index” measures contact intensity to co-workers and customers at work based on O*NET data at the five-digit industry level on a scale between 0 and 100 (see main text and footnote 21). “Work-from-Home Capacity” indicates the percentage share of workers in each two-digit industry able to work from home provided by Alipour, Falck, and Schüller (2021). “High Contact Intensity” is an indicator that is one if the industry is classified as “high contact goods and services” following a heuristic classification by Eichenbaum et al. (2020). The number of observations refers to the number of firms in the ifo-BCS wave of 2020:M4–2020:M5 and set to “NA” due to data protection issues if the number of observations per industry is smaller than or equal to three.

Table A.3. Effects of the COVID-19 Pandemic on Planned Price Adjustment

	Planned Price Decrease			Planned Price Increase			Planned Price Change		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Strongly Neg.	0.16*** (0.0075)	0.11*** (0.0090)	0.10*** (0.0099)	-0.13*** (0.0056)	-0.034*** (0.0093)	-0.041*** (0.0099)	0.032*** (0.0086)	0.074*** (0.012)	0.064*** (0.013)
Negative	0.087*** (0.0074)	0.050*** (0.0085)	0.051*** (0.0087)	-0.14*** (0.0059)	-0.022*** (0.0088)	-0.023*** (0.0089)	-0.057*** (0.0089)	0.028*** (0.012)	0.028*** (0.012)
Positive	-0.0092 (0.0092)	-0.023*** (0.011)	-0.022* (0.012)	0.033 (0.021)	0.078*** (0.021)	0.076*** (0.022)	0.023 (0.021)	0.054*** (0.023)	0.054*** (0.024)
Observations	145,856	124,068	123,904	145,856	124,068	123,904	145,856	124,068	123,904
Controls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Time FE	No	Yes	No	No	Yes	No	No	Yes	No
Time × Industry FE	No	No	Yes	No	No	Yes	No	No	Yes
Firm FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes

Note: This table reports estimates from linear regressions of indicators for planned price decreases/increases/changes on indicators for *COVID-19 impact* categories, based on Equation (3). *COVID-19 impact* measures the impact of COVID-19 on the current business situation on a seven-point scale from -3 (“negative”) to +3 (“positive”) in the ifo-BCS, which we group and label “Strongly Negative” (-3), “Negative” (-2), and “Positive” (+2 and +3), and the base category “Weak/No Impact” (-1, 0, or +1). Control variables include separate indicators for positive and negative responses to the questions about business situation, business expectations, and orders, all lagged by three months. Appendix C provides translations of all corresponding survey questions. Industry fixed effects are at the two-digit WZ08 level. Time fixed effects are at the month-year level. Standard errors in parentheses are clustered at the firm level. Sample: 2018:M1–2020:M5. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.4. Effects of the COVID-19 Pandemic on Realized Price Adjustment

	Realized Price Decrease			Realized Price Increase			Realized Price Change		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Strongly Neg.	0.089*** (0.0071)	0.046*** (0.0087)	0.042*** (0.0098)	-0.095*** (0.0044)	-0.020*** (0.0083)	-0.020*** (0.0090)	-0.0066 (0.0080)	0.026** (0.011)	0.023* (0.012)
Negative	0.044*** (0.0071)	0.024*** (0.0083)	0.027*** (0.0087)	-0.093*** (0.0049)	-0.013 (0.0079)	-0.015* (0.0081)	-0.049*** (0.0082)	0.012 (0.011)	0.012 (0.011)
Positive	-0.0019 (0.012)	-0.011 (0.014)	-0.016 (0.014)	0.043** (0.020)	0.024 (0.022)	0.035 (0.022)	0.041* (0.023)	0.013 (0.024)	0.019 (0.025)
Observations	122,440	106,794	106,657	122,440	106,794	106,657	122,440	106,794	106,657
Controls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Time FE	No	Yes	No	No	Yes	No	No	Yes	No
Time ×	No	No	Yes	No	No	Yes	No	No	Yes
Industry FE									
Firm FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes

Note: This table reports estimates from linear regressions of indicators for realized price decreases/increases/changes on indicators for *COVID-19 impact* categories, based on Equation (3). *COVID-19 impact* measures the impact of COVID-19 on the current business situation on a seven-point scale from -3 (“negative”) to +3 (“positive”) in the ifo-BCS, which we group and label “Strongly Negative” (-3), “Negative” (-2), and “Positive” (+2 and +3), and the base category “Weak/No Impact” (-1, 0, or +1). Control variables include separate indicators for positive and negative responses to the questions about business situation, business expectations, and orders, all lagged by three months. Appendix C provides translations of all corresponding survey questions. Industry fixed effects are at the two-digit WZ08 level. Time fixed effects are at the month-year level. Standard errors in parentheses are clustered at the firm level. Sample: 2018:M1–2020:M5. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5. Effects of the COVID-19 Pandemic on Planned Price Adjustment: Only Open Firms

	Planned Price Decrease			Planned Price Increase			Planned Price Change		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Strongly Neg.	0.16*** (0.0087)	0.10*** (0.011)	0.10*** (0.012)	-0.14*** (0.0064)	-0.034*** (0.011)	-0.032*** (0.011)	0.022** (0.010)	0.070*** (0.015)	0.070*** (0.015)
Negative	0.084*** (0.0077)	0.042*** (0.0094)	0.043*** (0.0097)	-0.14*** (0.0063)	-0.024** (0.0093)	-0.019** (0.0094)	-0.057*** (0.0095)	0.019 (0.013)	0.023* (0.013)
Positive	-0.00053 (0.0096)	-0.018 (0.012)	-0.0086 (0.013)	0.034 (0.021)	0.068*** (0.022)	0.063*** (0.023)	0.033 (0.022)	0.050** (0.024)	0.055** (0.025)
Observations	110,952	93,967	93,772	110,952	93,967	93,772	110,952	93,967	93,772
Controls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Time FE	No	Yes	No	No	Yes	No	No	Yes	No
Time ×	No	No	Yes	No	No	Yes	No	No	Yes
Industry FE									
Firm FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes

Note: This table reports estimates from linear regressions of indicators for planned price decreases/increases/changes on indicators for *COVID-19 impact* categories, based on Equation (3). *COVID-19 impact* measures the impact of COVID-19 on the current business situation on a seven-point scale from -3 (“negative”) to +3 (“positive”) in the ifo-BCS, which we group and label “Strongly Negative” (-3), “Negative” (-2), and “Positive” (+2 and +3), and the base category “Weak/No Impact” (-1, 0, or +1). Control variables include separate indicators for positive and negative responses to the questions about business situation, business expectations, and orders, all lagged by three months. Appendix C provides translations of all corresponding survey questions. Industry fixed effects are at the two-digit WZ08 level. Time fixed effects are at the month-year level. Standard errors in parentheses are clustered at the firm level. Sample: 2018:M1–2020:M5, only firms that are not closed in response to the COVID-19 pandemic according to SQ4 of the ifo-BCS. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7. Source of Planned Price Increases across Sectors in July 2021

	Planned Price Increases							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Oil Exposure	0.32*** (0.082)				0.31*** (0.082)	0.062 (0.10)	-0.11 (0.11)	
Supply Shortages		0.26* (0.14)				0.27* (0.15)	0.46** (0.18)	
Labor Shortages			-0.31 (0.21)		-0.22 (0.19)	-0.071 (0.32)	(0.30)	
Energy Costs				0.0024 (0.018)			0.033 (0.022)	
Retail/Wholesale								0.34*** (0.034)
Manufacturing								0.28*** (0.032)
Constant	0.33*** (0.035)	0.40*** (0.083)	0.53*** (0.069)	0.51*** (0.046)	0.40*** (0.072)	0.38** (0.14)	0.15 (0.16)	0.24*** (0.023)
Observations	60	25	60	22	60	25	22	60
R ²	0.209	0.122	0.036	0.001	0.227	0.144	0.320	0.672
Manufacturing	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Retail/Wholesale	Yes	Yes	Yes	No	Yes	Yes	No	Yes
Services	Yes	No	Yes	No	Yes	No	No	Yes

Note: This table reports estimates from linear regressions of the share of planned price increases on goods supply shortage, skilled labor supply shortage, oil price exposure, and energy costs on two-digit NACE industry level in July 2021. Goods supply shortage is defined as the share of firms in a given sector that report to be constrained by lack of material/intermediates. Skilled labor supply shortage is defined as the share of firms in a given sector that report to be constrained by lack of skilled labor. Appendix C provides the corresponding translated survey question Q7. Oil price exposure is the share of oil and oil-related (measured as the input share of coking and oil products times the share of oil input in that sector) inputs in total production from 2016 input-output matrix. We winsorize the oil exposure share in four sectors at 1. Energy costs is the total share of energy in gross production value. Both statistics are obtained from the Federal Statistical Office. Goods supply shortages are only available for manufacturing and retail/wholesale sectors. Energy costs are only available for the manufacturing sectors. We weight regressions by the number of observations in each industry. Standard errors in parentheses are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix B. Descriptive Analysis of *COVID-19 Impact*

B.1 COVID-19 Impact in the ifo Business Climate Survey

This appendix provides an extensive descriptive analysis of *COVID-19 impact* based on the survey waves of April and May 2020, summarized in Section 3. Table B.1 shows summary statistics for each *COVID-19 impact* category, which we refer to as “strongly negative” (−3), “negative” (−2), “weakly negative” (−1), and “no impact” (0), with analogous labels for the positive categories.

First, the bottom panel documents substantial heterogeneity in *COVID-19 impact*. In April and May 2020, 33 percent of all firms report a strongly negative impact, 24 percent are negatively affected, while 11 percent experience no impact. On the other hand, a smaller share of in total 10 percent tells of at least weakly positive effects. The difference between the 90th and the 10th percentile amounts to 3 points in April 2020 (May 2020: 4 points) on the *COVID-19 impact* scale, and the same difference within two-digit industries equals 2.82 scale points (May 2020: 2.83 points). Hence, there is substantial firm-level heterogeneity in *COVID-19 impact*, mirroring the large within-industry variation of price decreases shown in Section 2.

Second, panel A shows that the share of firms that report positive business conditions increases monotonically with *COVID-19 impact*, and vice versa (bar one exception). A similar pattern emerges for business expectations. Note that business situation and business expectations are not sufficient to explain *COVID-19 impact*. Some firms with no or positive impact report contemporaneously negative business conditions and expectations, while others with no or negative impact report positive business conditions and expectations. Hence, *COVID-19 impact* captures independent information specific to the pandemic.

Third, manufacturing capacity utilization in April 2020 on average strongly decreased year-on-year in negatively affected firms, and increased in positively affected firms. For instance, manufacturing firms hit hardest operate at about 54 percent of potential.

Fourth, the April 2020 survey questionnaire asked firms about the expected percent change in revenues due to the *COVID-19* crisis. As panel A shows, this figure decreases with *COVID-19 impact*

Table B.1. Summary Statistics by COVID-19 Impact


	COVID-19 Impact							Total
	-3	-2	-1	0	1	2	3	
Panel A								
Positive Business Conditions	0.008	0.039	0.208	0.524	0.593	0.793	0.865	0.188
Negative Business Conditions	0.091	0.194	0.406	0.5	0.492	0.406	0.343	0.39
	0.907	0.545	0.138	0.038	0.029	0.029	0.059	0.467
	0.29	0.498	0.345	0.19	0.168	0.167	0.235	0.499
Positive Business Expectations	0.12	0.115	0.097	0.09	0.167	0.288	0.336	0.122
Negative Business Expectations	0.325	0.319	0.296	0.286	0.373	0.453	0.474	0.327
	0.743	0.673	0.541	0.3	0.252	0.217	0.229	0.581
	0.437	0.469	0.498	0.458	0.434	0.413	0.421	0.493
Capacity Utilization (Mfg., in %)	54.432	70.4	79.662	85	87.901	81.333	84.5	71.439
	20.279	18.1	15.404	13.676	12.244	21.674	19.527	20.864
Cap. Util. (Mfg., y-o-y, in pp)	-27.414	-13.911	-4.776	-0.222	4.538	0.132	11.176	-11.954
	21.945	16.839	14.354	13.885	10.335	18.177	11.254	20.222
Expected Revenue Change in %	-37.109	-22.922	-12.658	-6.232	-0.248	3.956	13.365	-20.973
	21.742	13.522	9.88	9.046	9.56	15.255	40.19	21.357
Panel B								
Contact Intensity Index (0-100)	58.611	56.31	55.011	54.832	57.11	57.933	58.911	56.758
	5.111	4.297	4.255	4.792	5.575	5.205	5.435	4.995
Work-from-Home (in %)	54.136	56.958	60.915	60.958	58.688	56.813	53.702	57.366
	12.27	13.056	14.163	14.191	13.491	12.339	11.597	13.466
High Contact Intensity (in %)	41.805	27.371	22.343	21.212	44.846	58.544	64.574	32.874
	49.331	44.595	41.663	40.897	49.779	49.343	47.936	46.978

(continued)

Table B.1. (Continued)

	<i>COVID-19 Impact</i>						Total
	-3	-2	-1	0	1	2	3
<i>Panel C</i>							
Planned Price Increase	0.089	0.076	0.095	0.107	0.182	0.212	0.312
	0.284	0.266	0.294	0.309	0.387	0.409	0.464
Planned Price Decrease	0.219	0.143	0.091	0.044	0.042	0.044	0.05
	0.414	0.35	0.288	0.205	0.201	0.206	0.219
Planned Price Change	0.308	0.219	0.187	0.151	0.224	0.256	0.362
	0.462	0.414	0.39	0.358	0.418	0.437	0.426
Price Increase	0.041	0.051	0.057	0.079	0.168	0.189	0.321
	0.199	0.219	0.232	0.27	0.374	0.392	0.468
Price Decrease	0.154	0.099	0.057	0.041	0.045	0.041	0.062
	0.361	0.298	0.232	0.199	0.208	0.199	0.242
Price Change	0.195	0.149	0.114	0.121	0.213	0.23	0.383
	0.396	0.356	0.318	0.326	0.41	0.422	0.487
Observations	3,626	2,596	2,400	1,254	553	317	223
Percent	33.06	23.67	21.88	11.43	5.04	2.89	2.03
Note: This table depicts means and standard deviations (smaller numbers below) by <i>COVID-19 impact</i> . Row variables in panels A and C are indicators for positive/negative business conditions/expectations or planned/realized price increases/decreases/changes, capacity utilization, year-on-year change in capacity utilization, and expected percentage changes in revenue due to COVID-19 from the ifo-BCS. <i>COVID-19 impact</i> measures the impact of COVID-19 on the current business situation on a seven-point scale from -3 ("negative") to +3 ("positive") in the ifo-BCS. Appendix C provides translation of all corresponding survey questions. Row variables in panel B are industry-specific proxy variables for risk of COVID-19 infection. "Contact Intensity (0-100)" measures contact intensity to co-workers and customers at work based on O*NET data at the five-digit industry level (see main text and footnote 21). "Work-from-Home (in %)" indicates the share of workers in each two-digit industry able to work from home provided by Alipour, Falck, and Schüller (2021). "High Contact Intensity (in %)" refers to the share of firms classified as "high contact goods and services" at the two-digit industry level following a heuristic classification by Eichenbaum et al. (2020). Table A.2 in Appendix A documents how we sort industries according to these three proxy variables for infection risk. Main sample: 2020:M4-2020:M5; capacity utilization in manufacturing and expected revenue change: 2020:M4.							
							10,969
							100

Figure B.1. Contact Intensity and *COVID-19 Impact* by Industry




Note: This figure displays the relationship between *COVID-19 impact* and the contact intensity index based on O*NET survey data described in footnote 21 of Section 3.2, averaged at the two-digit industry level. *COVID-19 impact* measures the impact of COVID-19 on the current business situation on a seven-point scale from -3 (“negative”) to $+3$ (“positive”) in the ifo-BCS in 2020:M4. Blue (green) dots represent industries above (below) the median of the work from home capacity index by Alipour, Falck, and Schüller (2021). The red line displays the linear fit, weighted by the number of firms per industry. Industries with less than 20 observations not shown. The full industry-level data underlying this figure (including omitted industries) is summarized in Table A.2 in Appendix A.

and strongly negatively affected firms on average expect a 37 percent fall.

Fifth, panel B relates *COVID-19 impact* to industry-specific proxy variables for risk of COVID-19 infection described in Section 3.2. Panel B of Table B.1 shows that the relationship between *COVID-19 impact* and these proxy variables is, if anything, weak and health risks are only slightly higher among firms reporting a strongly negative impact.

Figure B.1 plots mean *COVID-19 impact* against the mean O*NET contact intensity measure by two-digit sector. All industries on average report to be negatively affected by the COVID-19 pandemic. The hardest hit businesses are in services—for example, travel

Figure B.2. *COVID-19 Impact* and Adverse Supply and Demand Shifts



Note: This figure displays the relationship between proxy variables for adverse supply and demand shifts against the *COVID-19 impact* that measures the impact of COVID-19 on the current business situation on a seven-point scale from -3 (“negative”) to $+3$ (“positive”) reported to the ifo-BCS in 2020:M4 and 2020:M5. The green bars represent the share of firms for each value of *COVID-19 impact*. The solid blue line displays the share of firms that report supply shortages (asked in online panel of manufacturing and retail/wholesale industries in 2020:M4) over *COVID-19 impact* categories. The dashed orange line displays the share of firms that report low order backlog over *COVID-19 impact* categories (asked in 2020:M4 and 2020:M5). Appendix C provides translations of all corresponding survey questions.

arrangement and reservation services, hotels, and restaurants—which also tend to be relatively contact intensive. The retail industry is slightly more adversely affected than the wholesale industry. About 17 percent of all retailers report positive impact, reflecting mostly grocery stores. Most adversely affected sectors in manufacturing include beverage manufacturing, furniture, and cars. Overall, the relation between *COVID-19 impact* and contact intensity is negative at the industry level, although substantial variation remains holding contact intensity fixed. High-contact industries are also less able to substitute to work from home, as shown by the green dots.

Figure B.2 further correlates *COVID-19 impact* with proxy variables for adverse supply and demand shifts. First, the dashed orange

line shows that *COVID-19 impact* strongly correlates with the responses to a question about low order backlogs (Q3 in Appendix C). More than 85 percent of all strongly negatively affected firms report a low order backlog. Table B.2 shows that this pattern is most pronounced among services firms (93 percent), but also very widespread in the manufacturing (84 percent) and retail/wholesale industries (85 percent). Accordingly, more than half of all strongly positively affected firms report a high order backlog. Second, the solid blue line shows that *COVID-19 impact* also correlates with the responses to a question about supply chain disruptions (SQ2 in Appendix C). In April 2020, more than 40 percent of manufacturing firms in each negative category lacked intermediate products, and about 50 percent of retail and wholesale firms in each negative category reported supply shortages of final goods. These shares decrease as *COVID-19 impact* becomes weaker but increase again in positively affected firms, consistent with the notion of excess demand for the goods and services of these firms.

In sum, *COVID-19 impact* appears to capture shifts in economic activity related to both supply and demand forces and, to a lesser extent, health risks due to the COVID-19 crisis.

B.2 COVID-19 Impact and Price-Setting Behavior Early in the Pandemic

This section provides a detailed analysis of the relationship between *COVID-19 impact* and firms' price-setting behavior in the early phase in the pandemic. The main results are summarized in Section 3.

Panel C of Table B.1 documents the relationship between *COVID-19 impact* and price-setting behavior at the onset of the pandemic. About 16.4 percent of firms changed prices in March and April, where 9.6 percent of firms decreased and 6.8 percent of firms increased prices. Firms strongly affected by COVID-19, both negatively and positively, change their prices more frequently than mildly affected firms. The frequency of price increases rises with *COVID-19 impact*, while the frequency of price decreases falls. Specifically, of those firms that are strongly adversely affected only 4 percent increase their prices, while 32 percent of those strongly positively affected do so. Conversely, of those firms that are strongly adversely

Table B.2. Summary Statistics: Proxy Variables for Supply and Demand by COVID-19 Impact


	COVID-19 Impact						
	-3	-2	-1	0	1	2	3
Panel A: Manufacturing							
Lack of Intermediate Inputs	0.451	0.456	0.418	0.297	0.365	0.485	0.625
Business Closure	0.498	0.499	0.494	0.458	0.485	0.508	0.5
Low Order Backlog	0.255	0.127	0.046	0.007	0.025	0.037	0
High Order Backlog	0.436	0.333	0.21	0.084	0.158	0.189	0
	0.844	0.668	0.318	0.09	0.051	0.123	0.049
	0.363	0.471	0.466	0.287	0.221	0.331	0.218
	0.027	0.035	0.102	0.243	0.325	0.494	0.683
	0.162	0.183	0.302	0.43	0.47	0.503	0.471
							0.426
							0.495
							0.126
							0.332
							0.535
							0.099
							0.299
Panel B: Retail/Wholesale							
Lack of Final Good Supply	0.497	0.548	0.533	0.368	0.566	0.61	0.661
Business Closure	0.5	0.498	0.5	0.485	0.498	0.491	0.477
Low Order Backlog	0.443	0.237	0.174	0.135	0.124	0.114	0.139
High Order Backlog	0.497	0.426	0.379	0.343	0.331	0.319	0.347
	0.851	0.625	0.387	0.174	0.133	0.093	0.141
	0.357	0.484	0.487	0.379	0.34	0.291	0.349
	0.024	0.045	0.105	0.151	0.266	0.399	0.592
	0.153	0.208	0.307	0.359	0.443	0.491	0.493
							0.523
							0.5
							0.268
							0.443
							0.533
							0.499
							0.119
							0.324

(continued)

Table B.2. (Continued)

COVID-19 Impact							
	-3	-2	-1	0	1	2	3
Panel C: Services							
Business Closure	0.244 0.43	0.047 0.211	0.036 0.187	0.028 0.166	0.027 0.164	0.02 0.14	0 0
Low Order Backlog	0.929	0.624	0.179	0.054	0.021	0.039	0.314
High Order Backlog	0.256	0.485	0.383	0.227	0.144	0.196	0.471
	0.015	0.042	0.179	0.365	0.448	0.647	0.543
	0.122	0.201	0.383	0.482	0.499	0.483	0.505
Note: This table depicts means and standard deviations (smaller numbers below) by <i>COVID-19 impact</i> and industry. Row variables are indicators for lack of intermediate inputs/final good supply, business closure, and high/low order backlog from the ifo-BCS. In retail/wholesale orders are assessed relative to the previous year. <i>COVID-19 impact</i> measures the impact of COVID-19 on the current business situation on a seven-point scale from -3 ("negative") to +3 ("positive") in the ifo-BCS. Appendix C provides translation of all corresponding survey questions. Sample: 2020:M4 for lack of supply measures, and 2020:M4–2020:M5 for business closure and orders.							

Figure B.3. COVID-19 Impact and Planned Price Adjustment by Industry




Note: This figure displays the relationship between the mean *COVID-19 impact* in 2020:M4 and the change in the frequency of planned price decreases (left) and planned price increases (right) between 2019:M4 and 2020:M4 at the two-digit NACE industry level. Blue (green) dots represent industries in which more (less) than 50 percent of firms report low orders. The red line displays the linear fit, weighted by the number of firms per industry. Industries with less than 20 observations not shown.

affected about 15 percent decrease their prices, while this share is only 6 percent among strongly positively affected firms. Looking ahead, about 23.8 percent of firms plan to change prices in the following three months, where 10.2 percent of firms plan to increase and 13.6 percent of firms plan to decrease prices. Hence, firms overall tend to decrease prices more often at the onset of the COVID-19 pandemic and also plan to decrease prices more frequently going forward. These price-setting patterns are prevalent in all sectors, in particular in retail/wholesale.

Figure B.3 plots the change in the frequency of planned price decreases and increases (relative to 12 months before to control for seasonality) against *COVID-19 impact* in April 2020 across industries. Blue (green) dots indicate industries in which more (less) than 50 percent of firms report low orders. The frequency of planned price decreases rises in most industries, especially so in industries with low demand. The frequency of planned price increases falls in most industries, especially so in industries with low demand. This exercise supports our interpretation that demand effects are important for firms adversely affected by COVID-19.

Figure B.4 plots the yearly change in the frequency of planned price adjustments in April 2020 against contact intensity across

Figure B.4. Contact Intensity and Planned Price Adjustment by Industry



Note: This figure displays the relationship between the mean contact intensity index based on O*NET survey data described in footnote 21 and the yearly change in the frequency of planned price decreases (left) and planned price increases (right) in 2020:M4 at the two-digit NACE industry level. Blue (green) dots represent industries above (below) the median of the work from home capacity index by Alipour, Falck, and Schüller (2021). The red line displays the linear fit, weighted by the number of firms per industry. Industries with less than 20 observations not shown.

industries. The frequency of planned price decreases, shown in panel A, rises in the bulk of industries. This increase is stronger in contact-intensive businesses such as hotels, restaurants, and retail, suggesting a relationship between health risks and planned price decreases. Holding contact intensity fixed, the remaining variation in the frequency of planned price decreases remains substantial and economically large, however. Indeed, the R^2 for the linear fit shown is 0.16. Thus, health risks appear to only explain parts of the observed price adjustment patterns during the COVID-19 pandemic, while other factors unrelated to health risks seem to be quantitatively more important.

Similarly, the frequency of planned price increases falls in all but three industries. Panel B of Figure B.4 shows that only paper products (including the toilet paper industry), financial services, and education see an increase the frequency of planned price increases. Contact-intensive businesses tend to have a more pronounced although insignificant decrease, and the variation holding contact intensity fixed remains substantial.

Highly contact-intensive industries have been forced to close under the lockdown in the early phase of the pandemic. An advantage of our approach is that we observe planned price changes for businesses that are currently closed due to the COVID-19 pandemic (question SQ4 in our survey data). Price plans for closed firms refer to price changes that firms anticipate when reopening. Clearly, some sectors and industries are potentially more strongly affected by business closures than others, e.g., hotels and restaurants. Table B.3 compares planned price changes for open and closed businesses in April 2020. There are no stark differences, but closed firms expect to change prices more often than open firms. In retail, closed firms expect to decrease prices more often, which hints at a relatively higher importance of demand. Hotels and restaurants that are closed expect to increase prices more than open business, which suggests that supply shortages are relatively more important.

Table B.3. Summary Statistics: Planned Price Changes in Closed and Open Firms

	Full Sample		Manufacturing		Services		Retail		Hotel/Rest.	
	Closed	Open	Closed	Open	Closed	Open	Closed	Open	Closed	Open
Planned Price Increases	0.11	0.1	0.02	0.06	0.12	0.07	0.14	0.19	0.09	0.06
Planned Price Decreases	0.25	0.13	0.19	0.13	0.23	0.11	0.29	0.15	0.29	0.3
Number of Observations	866	5,151	214	1,924	201	1,895	451	1,332	106	130

Note: This table depicts means in indicators for planned price increases and decreases by sector and for hotels and restaurants.
Sample: 2020:M4.

Appendix C. Overview of Survey Questions

C.1 Regular Questions in the ifo-BCS

The following set of questions, which are asked regularly on a monthly basis in the ifo-BCS, are used in this paper (English translation of German original).

Q1: Planned Price Changes:

Plans and expectations for the next 3 months: The prices of our goods/service will most probably [1] increase, [0] stay the same, or $[-1]$ decrease.

Q2: Realized Price Changes:²⁷

Trends in previous month: Compared to the month before, the prices of our goods/services [1] increased, [0] stayed the same, or $[-1]$ decreased.

Q3: Order Backlog/Demand:²⁸

Current situation: We evaluate our total backlog of orders as [1] comparatively large, [0] sufficient, or $[-1]$ too small.

Q4: Current Business Situation:

Current situation: We evaluate our current business situation as [1] good, [0] satisfactory, or $[-1]$ bad.

Q5: Expected Business Situation:

Expectations for the next six months: our business situation will be [1] more favorable, [0] stay approximately the same, or $[-1]$ more unfavorable.

²⁷In the ifo-BCS covering manufacturing and services firms, Q2 on realized price changes has been asked in the online panel since 2020:M4 only.

²⁸In the ifo-BCS on the retail/wholesale industries, Q3 is related relative to the situation one year before.

Q6: Capital Utilization [Manufacturing only, quarterly frequency]:

The utilization of our capacities is currently (normal full utilization = 100%):

- ☐ 30% ☐ 40% . . . ☐ 70% ☐ 75% . . . ☐ 100%
☐ if more than 100%: ____%

Q7: Constraints to Production/Business Activity [quarterly frequency]:

Our domestic production/business activity is currently constrained:

- ☐ yes ☐ no

If yes, by the following factors:

- ☐ lack of demand
☐ lack of material/intermediates
☐ lack of skilled labor
☐ . . . [Multiple additional options, not used in this paper]

Q8: Credit Negotiations [quarterly frequency]:

We have conducted credit negotiations with banks in the past 3 months:

- ☐ yes ☐ no

If yes, the banks behaved:

- ☐ compliant
☐ normal
☐ restrictive

C.2 Supplemental Questions Related to the COVID-19 Crisis

The wording of the special questions related to the COVID crisis in the ifo-BCS were as follows:²⁹

²⁹Due to space limitations on the paper-based questionnaires, some questions were only asked in the online panel of the ifo-BCS that was used by more than 75 percent of the survey participants.

SQ1 COVID-19 Impact [asked in all waves between April 2020 and March 2021]:

Do you realize an effect of the Corona pandemic on your current business situation? Is this effect negative or positive?

negative ☐ -3 ☐ -2 ☐ -1 ☐ 0 ☐ +1 ☐ +2 ☐ +3 positive

SQ2 Shortage of Supply [asked in online panel of manufacturing and retail/wholesale industries in April 2020, only]:

[Manufacturing:] Are you currently affected by problems with a shortage of supply of important intermediate goods from within Germany or abroad?

[Retail/Wholesale:] Are you currently affected by problems with a shortage of supply of important goods from within Germany or abroad?

☐ Yes ☐ No

SQ3 Expected Revenue Change due to COVID Crisis [asked in April 2020]:

Which effect of the Corona pandemic on your turnover do you expect in the current year?

☐ Increase of ____% ☐ No effect ☐ Decline of ____%

SQ4 Business Closure [asked in April and May 2020]:

Which measures has your firm taken in response to the COVID-19 pandemic?

☐ Business closures³⁰

☐ . . . [Multiple additional options, not used in this paper]

³⁰Choices slightly differed between industry-specific surveys. Manufacturing: “plant closures/stop of production”; Retail/Wholesale: “Closure of sales/business outlets”; Services: “business closures.”

SQ5 Importance of Intermediates from Abroad [asked in April 2020; online panel of manufacturing industries only]:

- a) Did you rely on important shipments of goods from abroad before the Corona pandemic?
☐ Yes ☐ No
- b) If yes, did those important shipments originate from China, Italy, or any other heavily affected country?
☐ China ☐ Italy ☐ Other countries: _____

SQ6 Adverse Effect of COVID Crisis [asked June 2020 (online panel only, November 2020 (full sample), February 2021 (full sample), and June 2021 (online panel only)]:

Due to the COVID-19 crisis, we are currently experiencing adverse effects in the following areas:

- a) Finances (e.g., liquidity):
 No adverse effects ☐ ☐ ☐ ☐ ☐ Large adverse effects
- b) Domestic Market (e.g., demand, order situation):
 No adverse effects ☐ ☐ ☐ ☐ ☐ Large adverse effects
- c) Foreign Market (e.g., demand, order situation):
 No adverse effects ☐ ☐ ☐ ☐ ☐ Large adverse effects
- d) Personnel (e.g., absences, exemptions, shortage):
 No adverse effects ☐ ☐ ☐ ☐ ☐ Large adverse effects
- e) Purchasing (e.g., supply chains, warehousing):
 No adverse effects ☐ ☐ ☐ ☐ ☐ Large adverse effects
- f) Regulations by Government (e.g., closures, hygiene concepts):
 No adverse effects ☐ ☐ ☐ ☐ ☐ Large adverse effects

SQ7 [asked in March 2020]

If you experience negative effects due to the COVID-19 pandemic on your business, which are those?

- ☐ Declining demand
☐ Impairment of business operations of foreign subsidiaries
☐ Delay/cancellation of shipments of intermediate goods or raw materials

- ☐ Delay/cancellation of shipments of final goods
- ☐ Increasing prices for intermediate goods or raw materials
- ☐ Decline of production
- ☐ Necessity for increased stock-keeping
- ☐ Delay/cancellation of business trips
- ☐ Others, in particular: _____

References

- Addressi, W., M. Pulina, and F. Sallusti. 2017. "Impact of Changes in Consumer Preferences on Sectoral Labour Reallocation: Evidence from the Italian Economy." *Oxford Bulletin of Economics and Statistics* 79 (3): 348–65.
- Alekseev, G., S. Amer, M. Gopal, T. Kuchler, J. Schneider, J. Stroebel, and N. Wernerfelt. 2023. "The Effects of COVID-19 on U.S. Small Businesses: Evidence from Owners, Managers, and Employees." *Management Science* 69 (1): 7–24.
- Alipour, J.-V., O. Falck, and S. Schüller. 2021. "Germany's Capacities to Work from Home." CESifo Working Paper No. 8227 (updated version, November).
- Alvarez, S. E., and S. M. Lein. 2020. "Tracking Inflation on a Daily Basis." *Swiss Journal of Economics and Statistics* 156 (1): Article 18.
- Bachmann, R., B. Born, S. Elstner, and C. Grimme. 2019. "Time-Varying Business Volatility and the Price Setting of Firms." *Journal of Monetary Economics* 101 (January): 82–99.
- Baker, S. R., R. Farrokhnia, S. Meyer, M. Pagel, and C. Yannelis. 2020. "How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic." *Review of Asset Pricing Studies* 10 (4): 834–62.
- Balduzzi, P., E. Brancati, M. Brianti, and F. Schiantarelli. 2020. "The Economic Effects of COVID-19 and Credit Constraints: Evidence from Italian Firms' Expectations and Plans." Discussion Paper No. 13629, Institute of Labor Economics.
- Balleer, A., N. Hristov, and D. Menno. 2020. "Menu Costs, the Price Gap Distribution and Monetary Non-neutrality: The Role of Financial Constraints." CEPR Discussion Paper No. 11790.
- Balleer, A., and P. Zorn. 2019. "Monetary Policy, Price Setting, and Credit Constraints." CEPR Discussion Paper No. 14163.

- Baqae, D. R., and E. Farhi. 2022. "Supply and Demand in Disaggregated Keynesian Economies with and Application to the Covid-19 Crisis." *American Economic Review* 112 (5): 1397–1436.
- Barrero, J. M., N. Bloom, S. J. Davis, and B. H. Meyer. 2021. "COVID-19 Is a Persistent Reallocation Shock." *AEA Papers and Proceedings* 111 (May): 287–91.
- Bartik, A. W., M. Bertrand, Z. Cullen, E. L. Glaeser, M. Luca, and C. Stanton. 2020. "The Impact of COVID-19 on Small Business Outcomes and Expectations." *Proceedings of the National Academy of Sciences* 117 (30): 17656–66.
- Bloom, N., R. S. Fletcher, and E. Yeh. 2021. "The Impact of COVID-19 on US Firms." Working Paper No. 28314, National Bureau of Economic Research.
- Brinca, P., J. B. Duarte, and M. F. E. Castro. 2021. "Measuring Sectoral Supply and Demand Shocks during COVID-19." *European Economic Review* 139 (October): Article 103901.
- Buchheim, L., J. Dovern, C. Krolage, and S. Link. 2022. "Sentiment and Firm Behavior during the COVID-19 Pandemic." *Journal of Economic Behavior and Organization* 195 (March): 186–98.
- Buchheim, L., C. Krolage, and S. Link. 2022. "Sudden Stop: When Did Firms Anticipate the Potential Consequences of COVID-19?" *German Economic Review* 23 (1): 79–119.
- Caballero, R. J., and A. Simsek. 2021. "A Model of Endogenous Risk Intolerance and LSAPs: Asset Prices and Aggregate Demand in a 'Covid-19' Shock." *Review of Financial Studies* 34 (11): 5522–80.
- Cabral, L., and L. Xu. 2021. "Seller Reputation and Price Gouging: Evidence from the COVID-19 Pandemic." *Economic Inquiry* 59 (3): 867–79.
- Carvalho, V., J. R. García, S. Hansen, Á. Ortiz, T. Rodrigo, S. R. Mora, and P. Ruiz. 2021. "Tracking the Covid-19 Crisis with High-Resolution Transaction Data." *Royal Society Open Science* 8 (8): Article 210218.
- Cavallo, A. 2020. "Inflation with Covid Consumption Baskets." Working Paper No. 27352, National Bureau of Economic Research.
- Chetty, R., J. N. Friedman, N. Hendren, M. Stepner, and The Opportunity Insights Team. Forthcoming. "The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data." *Quarterly Journal of Economics*.

- Eichenbaum, M., S. Rebelo, and M. Trabandt. 2021. "The Macroeconomics of Epidemics." *Review of Financial Studies* 34 (11): 5149–87.
- Eichenbaum, M. S., M. G. De Matos, F. Lima, S. Rebelo, and M. Trabandt. 2020. "How Do People Respond to Small Probability Events with Large, Negative Consequences?" Working Paper No. 27988, National Bureau of Economic Research.
- Enders, Z., F. Hünnekes, and G. J. Müller. 2019. "Monetary Policy Announcements and Expectations: Evidence from German Firms." *Journal of Monetary Economics* 108 (December): 45–63.
- Fetzer, T. 2022. "Subsidizing the Spread of COVID-19: Evidence from the UK's Eat-Out-to-Help-Out Scheme." *Economic Journal* 132 (643): 1200–1217.
- Ganglmair, B., A. Kann, and I. Tsanko. 2021. "Markups for Consumers." *Journal of Economics and Statistics* 241 (5/6): 701–34.
- Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek. 2017. "Inflation Dynamics during the Financial Crisis." *American Economic Review* 107 (3): 785–823.
- Guerrieri, V., G. Lorenzoni, L. Straub, and I. Werning. 2022. "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" *American Economic Review* 112 (5): 1437–74.
- Hassan, T. A., S. Hollander, L. Van Lent, and A. Tahoun. Forthcoming. "Firm-level Exposure to Epidemic Diseases: Covid-19, SARS, and H1N1." *Review of Financial Studies*.
- Jaravel, X., and M. O'Connell. 2020. "Real-time Price Indices: Inflation Spike and Falling Product Variety during the Great Lockdown." *Journal of Public Economics* 191 (November): Article 104270.
- Kaplan, G., B. Moll, and G. L. Violante. 2020. "The Great Lockdown and the Big Stimulus: Tracing the Pandemic Possibility Frontier for the U.S." Working Paper No. 27794, National Bureau of Economic Research, Inc.
- Kim, R. 2020. "The Effect of the Credit Crunch on Output Price Dynamics: The Corporate Inventory and Liquidity Management Channel." *Quarterly Journal of Economics* 136 (2): 563–619.
- Lehmann, R. 2023. "The Forecasting Power of the ifo Business Survey." *Journal of Business Cycle Research* 19 (1): 43–94.

- Lein, S. M. 2010. "When Do Firms Adjust Prices? Evidence from Micro Panel Data." *Journal of Monetary Economics* 57 (6): 696–715.
- Link, S. 2019. "The Price and Employment Response of Firms to the Introduction of Minimum Wages." CESifo Working Paper No. 7575.
- . 2020. "Harmonization of the ifo Business Survey's Micro Data." *Journal of Economics and Statistics* 240 (4): 543–55.
- Meyer, B. H., B. Prescott, and X. S. Sheng. 2022. "The Impact of the COVID-19 Pandemic on Business Expectations." *International Journal of Forecasting* 38 (2): 529–44.
- Mongey, S., L. Pilossoph, and A. Weinberg. 2021. "Which Workers Bear the Burden of Social Distancing?" *Journal of Economic Inequality* 19 (3): 509–26.
- Nakamura, E., and J. Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123 (4): 1415–64.
- . 2010. "Monetary Non-neutrality in a Multisector Menu Cost Model." *Quarterly Journal of Economics* 125 (3): 961–1013.
- Sauer, S., and K. Wohlrabe. 2019. "CEO or Intern—Who Actually Answers the Questionnaires in the ifo Business Survey?" *CESifo Forum* 20 (2): 29–31.
- . 2020. "ifo Handbuch der Konjunkturumfragen." ("ifo Handbook of Business Surveys.") ifo Beiträge zur Wirtschaftsforschung (ifo Contributions to Economic Research) No. 88.
- Yagan, D. 2015. "Capital Tax Reform and the Real Economy: The Effects of the 2003 Dividend Tax Cut." *American Economic Review* 105 (12): 3531–63.

How Wages Respond to the Job-Finding and Job-to-Job Transition Rates: Evidence from New Zealand Administrative Data*

Christopher Ball,^a Nicolas Groshenny,^{b,e} Özer Karagedikli,^{c,e,f}
Murat Özbilgin,^d and Finn Robinson^a

^aReserve Bank of New Zealand

^bLe Mans Université, GAINS-TEPP (CNRS 2042)

^cAsia School of Business

^dNew Zealand Treasury

^eCentre for Applied Macroeconomic Analysis

^fUniversity of Marburg

We use administrative data from New Zealand and exploit regional variations to evaluate the predictive power for wage dynamics of the job-finding and job-to-job transition rates. We find that the job-finding rate from unemployment plays a role in describing the wage dynamics of newly hired workers even after controlling for the job-to-job transition rate. The wages of new hires are much more responsive to both transition rates than the wages of job stayers. We then distinguish between the new hires transitioning from employment (job switchers) and the new hires coming from unemployment. The wages of job switchers are primarily related to the pace of job-to-job reallocation, and less significantly, to the job-finding rate. The wages of new hires from unemployment are exclusively linked to the job-finding rate and this association is stronger at the lower half of the wage distribution. Additionally, the wages of new hires from unemployment are more responsive to the job-finding rate than the wages of job stayers. The job-to-job transition rate has no impact on the wage dynamics of job stayers once the job-finding rate and the transition rate from inactivity to employment are controlled for.

JEL Codes: J31, J64.

*The authors would like to thank two anonymous referees, Bernd Hayo, Dean Hyslop, George Kudrna, François Langot, Anella Munro, Adrian Pagan, Ole Rummel, Brian Silverstone, Anthony Terriau, and Giulio Zanella for their comments and suggestions, and seminar participants at the Asia School of Business, Reserve Bank of New Zealand, Le Mans University, the ACE 2019 Conference in Melbourne, and the NZAE 20190 Conference in Wellington for comments. Corresponding author: Özer Karagedikli, ozer.karagedikli@asb.edu.my.

1. Introduction

Search-and-matching models are widely used to analyze wage dynamics. In the canonical search-and-matching model of the labor market due to Diamond, Mortensen, and Pissarides (henceforth the DMP model), the pace at which unemployed workers find jobs—the job-finding rate—is a crucial factor determining cyclical wage fluctuations. In that model, wages are set through bilateral bargaining between the employer and the employee. The bargaining power of a worker is determined by the attractiveness of their outside option, namely joining the pool of unemployed workers to look for another job. In the case of Nash bargaining, the equilibrium wage is a weighted average of a worker's productivity and her reservation wage, where the latter is directly influenced by the job-finding rate (the share of unemployed workers who transition to employment in a given period). According to the DMP model, when the job-finding rate is high, workers have more bargaining power so that, all else equal, wages increase.¹

A key assumption in the DMP model is that unemployed workers are the only source of labor for firms to fill their vacancies. Put differently, an employed person has to first become unemployed before she can start seeking another job. This assumption ignores two groups of workers: employed people who are searching for jobs (on-the-job search) and people who are not classified as part of the labor force but who may nonetheless be willing to work and are perhaps casually looking for jobs.

Moscarini and Postel-Vinay (2016) (henceforth MPV) observe that, in the Burdett and Mortensen (1998) model (henceforth the BM model), the job-to-job transition rate is the primary driver of wages. Competition between firms for workers who are already employed drives real wages higher through two channels: a strategic effect which benefits both job stayers and job movers, and a composition effect which only benefits job movers. In the BM model, the job-finding rate plays no role in shaping wage dynamics. This prediction differs strikingly from the DMP model. Motivated by this insight, MPV analyze U.S. aggregate time-series data on the

¹See Pissarides (2000) or Petrosky-Nadeau and Wasmer (2017) for a textbook treatment of the model.

job-finding rate and the job-switching rate.² MPV find that the evolution of wages over the business cycle is closely linked to the pace of job-to-job transitions and less so to variations in the job-finding rate. They interpret their findings as empirical support in favor of the BM model against the DMP model.³

Karahan et al. (2017) further assess the relative explanatory power of these two views of wage setting. They employ a panel data set that exploits state-level variations to measure the relative influence of the job-finding and job-to-job transition rates on cyclical wage fluctuations. They find that the wage dynamics of new hires and job stayers are both tightly linked to the pace of job-to-job transitions. Moreover, the explanatory power of the job-finding rate vanishes once they control for job-to-job flows. Their findings thus support the view that on-the-job search is the prevailing factor behind wage dynamics in the United States.⁴

Our paper is closely related to Karahan et al. (2017). We use New Zealand administrative data from the Linked Employer-Employee Data (LEED) and exploit regional variations to assess the explanatory power of the job-finding and job-to-job transition rates for wage fluctuations. We find that new hire earnings are tightly linked to the pace of job-to-job reallocation: A 1 percentage point increase in the job-to-job transition rate yields a 1.61 percent rise in the earnings of newly hired workers. However, contrary to Moscarini and Postel-Vinay (2016) and Karahan et al. (2017), we find that the job-finding rate from unemployment plays a role in describing the wage dynamics of newly hired workers even after controlling for the job-to-job transition rate. The estimated semi-elasticity of earnings to the job-finding rate is 0.30 (significant at the 5 percent level) for new hires. We also find that the wages of job stayers are much less

²The aggregate data were constructed by Fallick and Fleischman (2004) using the CPS.

³One should note that it is the present discounted value of wages (the user cost of labor) that matters for firms' hiring decision in the DMP model. Kudlyak (2014) and Basu and House (2016) show that the user cost of labor is even more procyclical than what wage series for new hires suggest.

⁴Fallick and Fleischman (2004), Faberman and Justiniano (2015), and Mukoyama, Patterson, and Şahin (2018) provide further evidence on the importance of on-the-job search in the United States. For related evidence from New Zealand and Australia, see Karagedikli (2018) and Deutscher (2019), respectively.

responsive to the labor market conditions than the wages of newly hired workers. For job stayers, the semi-elasticities of earnings to the job-finding and job-to-job transition rates are, respectively, 0.07 and 0.22 (both significant at the 5 percent level).

We then investigate whether our baseline results remain robust when we control for the rate of transitions from inactivity to employment (the *NE* rate). Once we control for *NE* flows, the explanatory power of the job-to-job transition rate becomes insignificant for stable earnings. Thus, the strategic channel of the BM model might not be relevant in New Zealand. Job-stayer earnings then appear to be connected only to the job-finding rates from unemployment and from inactivity. This result suggests that the cyclicity of workers' attachment to the labor market, a feature overlooked in the BM and DMP models, could matter for the dynamics of stable earnings. This seems consistent with existing empirical evidence for the United States.⁵ However, the LEED data do not allow us to disentangle migrants who start working quickly after their arrival in New Zealand from *NE* flows. Hence, we prefer to interpret the influence of *NE* transitions on stable earnings with a pinch of salt.

Turning to new hires, we find that controlling for the *NE* transition rate has no impact on our baseline results: the coefficients of the job-finding and job-to-job rates remain unchanged, 0.29 and 1.56, respectively (both highly significant). The coefficient of the *NE* rate is imprecisely estimated and barely significant.

Finally, to discover the source of the influence of the job-finding rate on new-hire wages, we focus on the provenance of newly hired workers. In doing so, we connect with the broader debate on the lack of wage rigidity in the data. Gertler, Huckfeldt, and Trigari (2020) argue that the pronounced procyclicality of new-hire wages observed in the data is due to workers switching jobs and does not apply to new hires coming from unemployment. The LEED data allow us to distinguish between new hires from employment (job switchers) versus new hires from unemployment. We then evaluate the predictive power of the job-finding and job-to-job transition rates for

⁵Moscarini and Postel-Vinay (2017) document a strong association between wages and the pace of *NE* flows. Elsby, Hobijn, and Şahin (2015) show that the participation margin accounts for one-third of unemployment volatility, while Armstrong and Karagedikli (2017) argue that the contribution of the participation margin could be even greater in New Zealand.

each group. We find the following results: (i) The earnings of new hires from unemployment are *exclusively* linked to the job-finding rate from unemployment, and this association is tighter at the bottom of the wage distribution. (ii) Contrary to Gertler, Huckfeldt, and Trigari (2020), the earnings of new hires from unemployment are more responsive to the labor market conditions (the job-finding rate from unemployment) than the earnings of job stayers. (iii) Consistent with a job-ladder mechanism (and with the composition effect in the BM model), the dominant predictor of job-switchers' earnings is the job-to-job transition rate, but the job-finding rate from unemployment retains some influence in the lower half of the distribution (the *NE* rate plays no role). (iv) The semi-elasticity of job-switchers' earnings to job-to-job flows is largest (equal to 3.43) at the lowest decile of the earnings distribution, highlighting the essential role of the bottom rung of the job ladder. This finding is consistent with recent evidence by Haltiwanger et al. (2018) for the United States.

2. Econometric Strategy

We first introduce the administrative data. Then we outline the empirical specifications used in our study.

2.1 Data

Our data come from a single source: the administrative Linked Employer and Employee Data (LEED) from the Inland Revenue Department (IRD).⁶ LEED covers the entire population who paid some form of pay-as-you-earn (PAYE) income tax in New Zealand. We compute the number of job switchers (people switching jobs within two consecutive quarters), job stayers (continuing in the same job), unemployed people entering employment (UE), and non-participants transitioning directly into employment (NE).⁷ We also calculate average nominal earnings for each of these groups.⁸

⁶The full data disclaimer is available on the last page of this paper.

⁷More specifically, LEED only enables us to classify as unemployed any person receiving an income support benefit, such as unemployment benefits.

⁸We use data on total earnings as in Karahan et al. (2017). We are unable to derive hourly earnings, as data on hours worked were not collected in LEED until April 2020.

Our quarterly observations use the middle month as the reference period, and are compared to the middle month of the previous quarter. In the New Zealand context, under the standard assumption that within-month job transitions are negligible, this approach should reduce the extent to which our quarterly data are affected by the time-aggregation bias.⁹ In addition, we use the highest income employer as a reference to reduce the number of spurious job transitions identified for those with multiple employers within a month.

We then construct the time series of nominal earnings, job-finding rates, and job-to-job transition rates for each of the 16 regions in New Zealand.¹⁰ Our data cover the 2001:Q1 to 2018:Q2 period. Figure 1 plots the job-to-job transition rate for each region. Before the global financial crisis, most regions had fluid labor markets with robust rates of job-to-job transition. Every region experienced a large decline in job-to-job flows around 2008–09, and several of them had not recovered to their pre-crisis levels by the end of our sample period.


2.2 Empirical Specifications

Our objective is similar to Karahan et al. (2017) where the empirical framework is not intended to establish a definitive causal relationship. Rather we want to distinguish between competing theories of labor market flows based on their respective prediction in terms of the association between wages, the job-finding rate, and the job-to-job transition rate.

⁹In terms of our quarterly data, we may be incorrectly assigning a job-to-job transition to someone who transitions to unemployment in the first or last month of the quarter while moving between employers in the observed middle months. However, there are stand-down periods in the New Zealand transfer system which result in a delay of about three weeks before new unemployment beneficiaries receive money (in place until 2020), which lower the likelihood of these transitions being miscoded. Moreover, employees need to give their current employers at least four weeks' notice before leaving their jobs. Hence the spurious job transition events outlined in Karahan et al. (2017) seem less likely to apply to our quarterly data.

¹⁰Karahan et al. (2017) also use nominal earnings. There are no regional deflator data available for New Zealand.

Figure 1. Regional Job-to-Job Transition Rates in New Zealand (seasonally adjusted)



Source: Statistics New Zealand, author estimates.

We use the microdata from New Zealand to empirically test the theoretical predictions of MPV about the relative explanatory power of the job-finding and job-to-job transition rates for wage growth. In line with Karahan et al. (2017), we estimate the following two fixed-effect specifications:

$$\log W_{it} = \alpha_i + \alpha_t + \beta_i t + \alpha_u \Lambda_{it}^u + \epsilon_{it} \quad (1)$$

$$\log W_{it} = \alpha_i + \alpha_t + \beta_i t + \alpha_u \Lambda_{it}^u + \alpha_e \Lambda_{it}^e + \epsilon_{it}, \quad (2)$$

where W_{it} denotes nominal earnings in region i in calendar quarter t . Λ_{it}^u is the transition probability from unemployment to employment in region i in quarter t . We compute it as the number of benefit recipients in region i in $t-1$ who enter employment in the same region within the next quarter, divided by the total number of benefit recipients in region i in $t-1$.¹¹ Λ_{it}^e is the regional job-to-job transition probability. We compute it as the share of employed people who transition from one employer in quarter $t-1$ to another in quarter t , with no observed intervening spell of non-employment.¹² The parameter α_i captures regional fixed effects, while the term $\beta_i t$ allows for a region-specific time trend. The time fixed effect α_t controls for variation in aggregate inflation and productivity as well as other aggregate cyclical factors. We check the validity of the fixed effects by using Hausmann's specification test. The test rejects the null hypothesis that the regional fixed effects are uncorrelated with the regressors, which confirms that our fixed-effect specification is appropriate.

Estimating specifications (1) and (2) replicates the approach of Karahan et al. (2017) for New Zealand.¹³ We then include the

¹¹Not all unemployed are registered to receive an unemployment benefit. Therefore, our measure of the job-finding rate is likely to exclude the very short-term unemployed, which are the most likely to get a job. This may bias our estimates of the relative predictive power of the job-finding rate downward. Below, we will explore this issue by controlling for *NE* transitions.

¹²We ensure that a person received continuous earnings throughout two consecutive quarters. We then check whether that person received income from the same employer or not.

¹³Similarly to Karahan et al. (2017), a caveat of our approach to test the predictions of the BM model is that Λ_{it}^e and Λ_{it}^u measure the realized transition rates rather than the arrival rates of job offers to employed and unemployed workers respectively.

regional transition probability from non-participation to employment, Λ_{it}^n , to obtain the following augmented specification:

$$\log W_{it} = \alpha_i + \alpha_t + \beta_i t + \alpha_u \Lambda_{it}^u + \alpha_e \Lambda_{it}^e + \alpha_n \Lambda_{it}^n + \epsilon_{it}. \quad (3)$$

Adding the *NE* transition rate should help us to capture the very short-term unemployed (who may not register for unemployment insurance and would therefore be counted as *NE* instead of *UE* flows) and to appraise the influence of the participation margin (Elsby, Hobijn, and Şahin 2015, Armstrong and Karagedikli 2017).¹⁴

3. Results

We start with the replication of Karahan et al. (2017) for New Zealand. We then discuss the findings from the augmented specification that aims to account for the participation margin. Finally, we present results for the earnings distribution of job stayers, job switchers, and new hires from unemployment.

3.1 Replication of Karahan et al. (2017) for New Zealand

MPV point out that, in the BM model, wages and job-to-job transitions interact through two mechanisms: a compositional effect and a strategic rent-extraction channel. The former mechanism follows from job switchers climbing up the job ladder: workers transit between jobs when they receive (and accept) a higher wage offer. The latter effect reflects the competition between firms to retain their workforce when workers have more outside options: employees may extract a wage increase by generating an outside offer and asking their current employer to match it. In that case, the worker does not switch jobs but still gets a pay rise. Both channels favor the job-to-job transition rate over the job-finding rate from unemployment in terms of explaining the dynamics of new-hire earnings. In addition, the strategic effect also implies a positive relationship between job-to-job transitions and job-stayer earnings.

¹⁴It is worth noting that our LEED-based measure of *NE* flows may include recent migrants to New Zealand, which we cannot identify separately. If migrants are self-selected and there are differences in average skill sets between natives and migrants, the interpretation of the coefficient α_n might not be straightforward.

Table 1. Main Results: Replication of Karahan et al. (2017)

	Stable Earners		New Hires		All	
	(1)	(2)	(1)	(2)	(1)	(2)
Λ^u	0.077** (0.030)	0.067** (0.025)	0.373** (0.140)	0.300** (0.122)	0.252** (0.088)	0.204** (0.076)
Λ^e		0.222** (0.099)		1.611*** (0.304)		1.054*** (0.180)
Note: Number of observations: 1,168. Robust standard errors (clustered at regional level) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.						

The New Zealand micro-data allow us to distinguish between job stayers and new hires. Haefke, Sonntag, and van Rens (2013) show that the wages of newly hired workers are much more volatile and procyclical than the wages of job stayers. We follow Karahan et al. (2017) and estimate Equations (1) and (2) for these two groups. The results are reported in Table 1.¹⁵ The first two columns of Table 1 report the results for stable earnings. When we only include the job-finding rate, the semi-elasticity of stable earnings to the job-finding rate is about 0.08 and significant at the 5 percent level. A 1 percentage point increase in the job-finding rate is accompanied by a 0.08 percent increase in earnings for job stayers. Controlling for the job-to-job transition rate has little impact on the explanatory power of the job-finding rate which, in contrast to Karahan et al. (2017), remains significant at the 5 percent level. The association between stable earnings and the job-to-job transition probability is stronger, with an estimated semi-elasticity equal to 0.22 and significant at the 5 percent level, providing some support for the strategic rent-extraction channel.¹⁶

The next two columns present the results for new-hire earnings. When included on its own, the job-finding rate is a significant

¹⁵We use the terms “job stayers” and “stable earners” interchangeably to denote people who remain employed in the same job.

¹⁶We will see in the next section that this support vanishes when we control for transitions from inactivity to employment.

predictor of earnings growth for new hires with an estimated semi-elasticity equal to 0.37. In stark contrast to Moscarini and Postel-Vinay (2016) and Karahan et al. (2017), when we control for the job-to-job transition rate, we find that the explanatory power of the job-finding rate does not evaporate: it declines slightly to 0.30 but remains significant at the 5 percent level. We find a tight link between new-hire earnings and the job-to-job transitions, with a highly significant semi-elasticity equal to 1.61. Our results confirm that the wages of newly hired workers are much more responsive to labor market conditions than the wages of job stayers, as pointed out by Haefke, Sonntag, and van Rens (2013).

The last two columns show the results for all earnings, with no distinction between new hires and stable earners. On its own, the job-finding rate is significant. Once we include the job-to-job transition rate, the explanatory power of the job-finding rate diminishes slightly, from 0.25 to 0.20, but remains significant. The coefficient estimate of the job-to-job transition rate is about five times larger than the one of the job-finding rate. A 1 percentage point increase in the job-to-job transition rate yields roughly a 1 percent rise in average earnings. To sum things up, compared to Karahan et al. (2017), the key difference is that the job-finding rate from unemployment plays a role in describing wage dynamics of newly hired workers even after controlling for the job-to-job transition rate.

3.2 *Taking the Participation Margin into Account*

As discussed above, the LEED data only allow us to classify as unemployed any person receiving an income support benefit, such as unemployment insurance benefit. Some of the short-term unemployed may not register for the unemployment benefit, especially the individuals who are more likely to find a job quickly (because they search harder or are more employable). Our procedure will mistakenly count these transitions as *NE* instead of *UE* flows and our measure of the job-finding rate will miss these individuals. We think this could work against the job-finding rate, and so might artificially bias the explanatory power of the job-to-job rate upward.

Controlling for the *NE* transition rate may help account for the unregistered short-term unemployed and, more generally, may shed light on the influence of the participation margin. Moscarini and

Table 2. Controlling for *NE* Transitions

	Stable Earners		New Hires		All	
	(2)	(3)	(2)	(3)	(2)	(3)
Λ^u	0.067** (0.025)	0.060** (0.022)	0.300** (0.122)	0.292** (0.121)	0.204** (0.076)	0.197** (0.075)
Λ^e	0.222** (0.099)	0.176 (0.102)	1.611*** (0.304)	1.559*** (0.298)	1.054*** (0.180)	1.004*** (0.177)
Λ^n		1.441*** (0.302)		1.633* (0.753)		1.575*** (0.506)
Note: Number of observations: 1,168. Robust standard errors (clustered at regional level) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.						

Postel-Vinay (2017) find evidence of a tight link between wages and *NE* transitions in the United States. Elsby, Hobijn, and Şahin (2015) document that flows in and out of the labor force are responsible for up to a third of the cyclical volatility of the unemployment rate. Armstrong and Karagedikli (2017) argue that the significance of the participation margin might be even larger in New Zealand.

To explore these issues, we estimate specification (3), which adds the regional inactivity-to-employment transition rates. To assess the robustness of our baseline findings to controlling for *NE* transitions, Table 2 compares the results from specifications (2) and (3).

For stable earners, we find that controlling for *NE* transitions alters the estimation results in two ways. First, the influence of the job-finding rate remains stable and significant, while the predictive power of the job-to-job rate becomes insignificant. This casts doubt on the significance of the strategic rent-extraction effect in the New Zealand context. The stable earners group consists of employed people who do not engage in search activity and of employed people who engaged in on-the-job search but decided to stay in the same job. The rent-extraction channel only applies to the latter group, and it might be hard to detect evidence of it when the strategically searching job stayers are scarce compared to the non-searching job stayers. Second, the semi-elasticity of stable earnings to *NE* flows is large (equal to 1.44) and highly significant. Note, however, that the LEED data do not enable us to distinguish between new migrants

and non-participants. Hence, we should interpret this coefficient cautiously. Wages of job stayers are manifestly related to the job-finding rates from inactivity and unemployment, pointing towards a link between stable earnings and workers' participation decisions.

For new-hire earnings, controlling for *NE* transitions has no effect on the baseline findings: the coefficients of the job-finding and job-to-job rates remain unchanged, at 0.29 and 1.56, respectively (both highly significant). The coefficient of the *NE* rate is imprecisely estimated and barely significant.

Interestingly, the relationship between *NE* transitions and earnings is much stronger for job stayers than for new hires. This is consistent with the view that our measure of *NE* flows partly reflect inflows of migrants who start working immediately. Coleman and Karagedikli (2018) and Howard (2020) provide evidence indicating that positive net migration flows boost aggregate demand and put broad-based upward pressures on prices throughout the economy, including the labor market, thereby lifting the average wage.

3.3 Looking at the Earnings Distribution of Stable Earners, Job Switchers, and New Hires from Unemployment

The LEED data allow us to go beyond the usual dichotomy between stable earners and new hires. To uncover the origin of the influence of the job-finding rate on new-hire wages, we follow Gertler, Huckfeldt, and Trigari (2020) and distinguish between new hires coming from unemployment (*UE* new hires) versus new hires switching from one job to another.¹⁷ Furthermore, to understand which part of the earnings distribution contributes most to the explanatory power of the job-finding rate, we estimate specification (3) for each group at each decile of that group's earnings distribution. Table 3 shows the results for stable earners, while Table 4 focuses on the two categories of new hires.

Looking at Table 3, we see a positive relationship between stable earnings and the job-finding rate at all deciles of the earnings

¹⁷For completeness, we have also considered new hires transitioning from inactivity. However, none of the transition rates appeared to be linked with the earnings of that group. These results are available upon request.

Table 3. Regressions by Deciles for Stable Earnings

Decile	Job Stayers		
	Λ^u	Λ^e	Λ^n
10 th	0.131*	0.722	1.154*
20 th	0.081**	0.295*	1.169***
30 th	0.056**	0.222*	1.334***
40 th	0.041*	0.176	1.356***
50 th	0.045*	0.112	1.246***
60 th	0.061**	0.094	1.301***
70 th	0.061**	0.103	1.299***
80 th	0.059***	0.069	1.393***
90 th	0.061***	0.081	1.451***
Note: Number of observations: 1,168. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.			

Table 4. Regressions by Deciles for New-Hire Earnings

Decile	UE New Hires			Job Switchers		
	Λ^u	Λ^e	Λ^n	Λ^u	Λ^e	Λ^n
10 th	1.796***	1.299	−3.248	0.674**	3.430***	3.429
20 th	2.040***	1.772	−2.596	0.486***	2.748***	1.558
30 th	1.871***	1.165	−2.505	0.329***	2.439***	1.173
40 th	1.480***	1.152	−1.749	0.249***	2.458***	1.202
50 th	1.166***	1.197	−0.706	0.218***	2.496***	1.164
60 th	0.931**	1.011	−0.357	0.161**	2.511***	1.455
70 th	0.696**	0.918	0.334	0.096	2.495***	1.370
80 th	0.404	1.031*	0.411	0.056	2.293***	0.741
90 th	0.021	1.003*	0.371	0.007	1.661***	1.025
Note: Number of observations: 1,168. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.						

distribution. The link is quantitatively small but consistently significant. The predictive power of the job-to-job rate is insignificant at most deciles, casting doubt on the strategic effect of the BM model in the New Zealand context. Stable earnings are tightly linked with the *NE* rate. This is in line with Moscarini and Postel-Vinay (2017), who find a close relationship between wages and *NE* transitions in the United States, as well as with Armstrong and Karagedikli (2017)

who argue that *NE* flows are particularly large in New Zealand. However, as discussed above, LEED does not enable us to disentangle migrants from non-participants, and we cannot rule out the conjecture that migration flows are a prime determinant of cyclical variations in stable earnings in New Zealand. A deeper analysis of the roles of migrants and non-participants in New Zealand wage dynamics is an interesting avenue for future research.

Turning to Table 4, we see that the earnings of newly hired workers transitioning from unemployment are *exclusively* linked to the job-finding rate. This association is stronger at the bottom of the distribution: the semi-elasticity is equal to 1.9 and highly significant at the three lower deciles. It then declines progressively as we move to higher deciles. Contrary to Gertler, Huckfeldt, and Trigari (2020), the earnings of new hires from unemployment are more responsive to the labor market conditions (the job-finding rate from unemployment) than the earnings of job stayers. Interestingly, the influence of the job-finding rate is stronger on new-hire wages at lower deciles of the wage distribution, in line with the findings of Katz and Krueger (1999) that the wage Phillips curve is more vivid for low-wage workers.

The earnings of job switchers are primarily connected with the job-to-job transition rate. This association is tight and highly significant (at 1 percent) throughout the entire earnings distribution. The semi-elasticity of job-switcher earnings to the job-to-job rate is 3.43 at the lowest decile. It then declines gradually to reach 1.66 at the top decile. Surprisingly, the job-finding rate also plays a significant role at the bottom half of the earnings distribution of job switchers. The explanatory power of the *NE* rate is insignificant at all deciles.

The results in this section include a number of insights, some of which are new, into the cyclical wage dynamics of newly hired workers. In contrast to Moscarini and Postel-Vinay (2016) and Karahan et al. (2017), our results show that the job-finding rate plays an important role in describing the cyclical wage dynamics of newly hired workers even after controlling for the job-to-job and *NE* rates. Indeed, the job-finding rate from unemployment is the only transition rate that consistently displays a significant relationship with the earnings of *all* categories of workers: job stayers, new hires from unemployment, and new hires switching jobs. Instead, once we control for the job-finding and *NE* rates, the job-to-job transition rate

appears to be exclusively connected to the earnings of new hires switching jobs, suggestive of a pure composition effect. The fact that the semi-elasticity of job-switcher earnings to the job-to-job rate is largest at the lowest decile of the earnings distribution further highlights the importance of the bottom rung of the job-ladder, in line with Haltiwanger et al. (2018).

4. Conclusions

In this paper, motivated by recent findings of Moscarini and Postel-Vinay (2016) and Karahan et al. (2017), we have attempted to distinguish between two models of wage determination—one in which the job-finding probability of the unemployed plays a key role (the DMP model), and another in which it does not play any role for the benefit of the job-to-job transition rate (the BM model). To do so, we have used administrative data from New Zealand and replicated the empirical strategy of Karahan et al. (2017), exploiting pooled cross-regional variations to assess the explanatory power of the job-finding and job-to-job transition rate for the cyclical wage dynamics of stable earners and new hires.

Although some of our results lend support to the view that on-the-job search is a prominent factor of cyclical wage dynamics in New Zealand—specifically, new-hire earnings are tightly linked to the pace of job-to-job reallocation—some others differ starkly from the views of Moscarini and Postel-Vinay (2016) as well as Karahan et al. (2017). In particular, we find that the job-finding rate retains significant explanatory power for the earning dynamics of newly hired workers even when controlling for the pace of job-to-job transitions. This result poses a challenge to the BM model. When we control for *NE* flows, the explanatory power of the job-to-job transition rate disappears for stable earnings, suggesting that the strategic rent-extraction channel of the BM model is not relevant in the New Zealand context. Instead, the participation margin seems to matter for the evolution of stable earnings, an aspect neglected in the BM and DMP model.

To discover the source of the influence of the job-finding rate on the cyclical wage dynamics of new hires, we have then distinguished between the new hires coming from the unemployment pool and those transitioning from one job to another. For each category, we

have evaluated the comparative explanatory power of the job-finding rate at different deciles of the earnings distribution. Earnings of new hires from unemployment are exclusively linked to the job-finding rate, and this association is especially tight at the bottom of the earnings distribution. Moreover, the earnings of *UE* new hires are more responsive to the labor market conditions (the job-finding rate from unemployment) than the earnings of job stayers, in contrast to evidence reported by Gertler, Huckfeldt, and Trigari (2020). Surprisingly, the job-finding rate also plays a role in describing the earnings dynamics of job switchers at the lower half of the earnings distribution. However, job switchers' earnings are primarily linked to the job-to-job transition rate, in agreement with the composition channel (i.e., a job-ladder mechanism) in the BM model. Furthermore, consistent with Haltiwanger et al. (2018), the bottom rung of the job ladder is the most essential one.

Our findings can inform recurrent policy debates on cyclical wage dynamics. Understanding the mechanisms of labor market dynamics and wage fluctuations is paramount to any central bank pursuing a flexible inflation-targeting strategy. In a number of countries, including New Zealand, the recent experience before the COVID-19 pandemic, marked by low unemployment rates and subdued wage growth left many policymakers perplexed.¹⁸ Our results highlight the prominence of a composition effect, working through a job-ladder mechanism, for wage growth in New Zealand. Besides, our findings reveal that the job-finding rate is linked to the wage dynamics of both stable earners and new hires, especially those coming from the unemployment pool.

Finally, Kudlyak (2014) and Basu and House (2016) show that it is the present discounted value of wages that matters for firms' hiring decision in the DMP model. This is especially the case in the presence of implicit contracts whereby wages of newly hired workers respond more to the outside option than wages of existing workers (Beaudry and DiNardo 1991), as we find in the New Zealand data. An interesting avenue for future research would be to investigate how the user cost of labor (the effective price for new hires) responds to the outside option in New Zealand data.

¹⁸Recent papers, such as Jørgensen and Lansing (2022), argue for an important role of expectations in the behavior of the Phillips curve.

Appendix. Robustness Checks

In this appendix, we conduct a number of robustness checks for the “All Employees” category. Our first robustness check is motivated by the observation in Figure 1 that the regional job-to-job series may contain a structural break around 2008: job-to-job flows in all regions experienced a sharp fall in 2008 and, in many regions, they did not recover fully over the subsequent decade. In order to detect a possible break, we run the Bai-Perron break test. The test signals the presence of a structural break in 2008:Q2. Based on this information, we split the sample at that date and re-estimate specifications (1) to (3) over each sub-period. Table 5 displays the results. Several comments are in order. First, the influence of the job-finding rate appears to have declined over time: the point estimate of the semi-elasticity falls from 0.325 in the first period to 0.157 in the second period. Although small in magnitude, the explanatory power of the job-finding rate remains highly significant in both periods. Second, the effect of job-to-job flows seems to have become somewhat more pronounced over time: the point estimate of the semi-elasticity of earnings with respect to the job-to-job probability increases moderately from 0.878 pre-2008 to 1.175 post-2008. The explanatory power of the job-to-job transition rate is highly significant in both periods. Finally, the influence of the inactivity-to-employment transition rate is highly unstable across the two periods. This striking instability calls for some caution when interpreting the influence of *NE* flows on earnings dynamics in New Zealand. As pointed out in the main text, a caveat of our analysis is that we cannot distinguish between *NE* flows and immigrants to New Zealand who start working shortly after their arrival. We leave that issue for future research.

To further explore the nature of the structural break, we perform a second robustness check. Instead of splitting the sample period in 2008, we now introduce a dummy variable, *break*, that takes the value 0 before 2008:Q2 and 1 afterwards. The dummy variable enters in specification (2) and (3) as an interaction term with the job-finding rate Λ^u , the job-to-job transition rate Λ^e , and the inactivity-to-employment transition rate Λ^n . The results are reported in Table 6. For both specifications, the slope interaction terms are insignificant, suggesting that the explanatory power of each transition rate stays relatively constant across the two periods.

Table 5. Pre- and Post-Break Estimation

	Pre-2008				Post-2008		
Λ^u	0.392*** (0.134)	0.326** (0.131)	0.325** (0.129)	Λ^u	0.195*** (0.061)	0.168*** (0.050)	0.157*** (0.049)
Λ^e		0.880*** (0.184)	0.878*** (0.180)	Λ^e		1.226*** (0.303)	1.175*** (0.301)
Λ^n			0.039 (0.731)	Λ^n			2.691*** (0.530)
Note: Number of observations: 1,168. Robust standard errors (clustered at regional level) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.							

Table 6. Full Sample Estimates with Break Interactions

	(2)	(3)
Λ^u	0.261** (0.110)	0.247** (0.114)
Λ^e	1.032*** (0.167)	1.025*** (0.188)
Λ^n		0.907 (0.881)
Break	0.626*** (0.016)	0.600*** (0.024)
$\Lambda^u \times \text{Break}$	0.083 (0.060)	-0.078 (0.068)
$\Lambda^e \times \text{Break}$	-0.023 (0.200)	-0.123 (0.248)
$\Lambda^n \times \text{Break}$		1.489 (1.713)
Note: Number of observations: 1,168. Robust standard errors (clustered at regional level) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.		

The break appears to be significant, but that comes from the intercept, not from the slope coefficients. Interestingly, the influence of Λ^n is now insignificant. Altogether, the two robustness checks convey the impression that our findings regarding the respective explanatory power of the job-finding rate and the job-to-job transition rate are reasonably robust, and that we should interpret the influence of *NE* flows on earnings with a pinch of salt.

Table 7. Controlling for Lagged Earnings

	(2)		(3)	
	(a)	(b)	(d)	(d)
Λ^u	0.204** (0.076)	0.215** (0.082)	0.197** (0.075)	0.207** (0.081)
Λ^e	1.054*** (0.179)	1.037*** (0.187)	1.004*** (0.177)	0.998*** (0.182)
Λ^n			1.575*** (0.505)	1.530*** (0.486)
lw(−1)		0.105* (0.059)		0.096 (0.057)
Note: Number of observations: 1,168. Robust standard errors (clustered at regional level) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.				

Our last robustness check tackles a different issue. For the sake of comparability with Karahan et al. (2017), our baseline specifications did not include the lagged log wage. These regressions all featured time fixed effects which were expected to capture the persistent element in wages. We now check this conjecture by including the lagged log wage in specifications (2) and (3). Looking at Table 7, we see that the lagged log wage term turns out to be economically and statistically insignificant, and its inclusion does not noticeably affect the coefficients of the job-finding rate and job-to-job transition rate. This confirms the adequacy of the baseline specifications with time fixed effects.

References

Armstrong, J., and Ö. Karagedikli. 2017. “The Role of Non-participants in Labour Market Dynamics.” Analytical Notes Series No. 2017/01, Reserve Bank of New Zealand.

Basu, S., and C. House. 2016. “Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models.” In *Handbook of Macroeconomics*, Vol. 2A, ed. J. B. Taylor and H. Uhlig, 297–354. North-Holland.

Beaudry, P., and J. DiNardo. 1991. “The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data.” *Journal of Political Economy* 99 (4): 655–88.

- Burdett, K., and D. Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review* 39 (2): 257–73.
- Coleman, A., and Ö. Karagedikli. 2018. "Residential Construction and Population Growth in New Zealand: 1996–2016." Discussion Paper No. 2018/02, Reserve Bank of New Zealand.
- Deutscher, N. 2019. "Job-to-Job Transitions and the Wages of Australian Workers." Working Paper 2019-07, Australian Treasury.
- Elsby, M. W. L., B. Hobijn, and A. Şahin. 2015. "On the Importance of the Participation Margin for Labor Market Fluctuations." *Journal of Monetary Economics* 72 (May): 64–82.
- Faberman, J., and A. Justiniano. 2015. "Job Switching and Wage Growth." Chicago Fed Letter No. 337.
- Fallick, B. C., and C. A. Fleischman. 2004. "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows." Finance and Economics Discussion Series No. 2004-34, Board of Governors of the Federal Reserve System.
- Gertler, M., C. Huckfeldt, and A. Trigari. 2020. "Unemployment Fluctuations, Match Quality, and the Wage Cyclicalities of New Hires." *Review of Economic Studies* 87 (4): 1876–1914.
- Haefke, C., M. Sonntag, and T. van Rens. 2013. "Wage Rigidity and Job Creation." *Journal of Monetary Economics* 60 (8): 887–99.
- Haltiwanger, J. C., H. R. Hyatt, L. B. Kahn, and E. McEntarfer. 2018. "Cyclical Job Ladders by Firm Size and Firm Wage." *American Economic Journal: Macroeconomics* 10 (2): 52–85.
- Howard, G. 2020. "The Migration Accelerator: Labor Mobility, Housing, and Demand." *American Economic Journal: Macroeconomics* 12 (4, October): 147–79.
- Jørgensen, P. L., and K. J. Lansing. 2022. "Anchored Inflation Expectations and the Slope of the Phillips Curve." Working Paper No. 2019-27, Federal Reserve Bank of San Francisco.
- Karagedikli, Ö. 2018. "Job-to-Job Flows and Inflation: Evidence from Administrative Data in New Zealand." Analytical Notes Series No. 2018/09, Reserve Bank of New Zealand.
- Karahan, F., R. Michaels, B. Pugsley, A. Şahin, and R. Schuh. 2017. "Do Job-to-Job Transitions Drive Wage Fluctuations over the Business Cycle?" *American Economic Review* 107 (5): 353–57.

- Katz, L. F., and A. B. Krueger. 1999. "The High-Pressure U.S. Labor Market of the 1990s." *Brookings Papers on Economic Activity* 30 (1): 1–88.
- Kudlyak, M. 2014. "The Cyclicalities of the User Cost of Labor." *Journal of Monetary Economics* 68 (November): 53–67.
- Moscarini, G., and F. Postel-Vinay. 2016. "Wage Posting and Business Cycles." *American Economic Review* 106 (5): 208–13.
- . 2017. "The Relative Power of Employment-to-Employment Reallocation and Unemployment Exits in Predicting Wage Growth." *American Economic Review* 107 (5): 364–68.
- Mukoyama, T., C. Patterson, and A. Şahin. 2018. "Job Search Behavior over the Business Cycle." *American Economic Journal: Macroeconomics* 10 (1): 190–215.
- Petrosky-Nadeau, N., and E. Wasmer. 2017. *Labor, Credit, and Goods Markets: The Macroeconomics of Search and Unemployment*, Vol. 1. MIT Press.
- Pissarides, C. A. 2000. *Equilibrium Unemployment Theory*, 2nd Edition, Vol. 1. MIT Press.

Integrated Data Infrastructure Disclaimer: The results in this report are not official statistics; they have been created for research purposes from the Integrated Data Infrastructure (IDI) managed by Statistics New Zealand.

Access to the anonymized data used in this study was provided by Statistics NZ in accordance with security and confidentiality provisions of the Statistics Act 1975. Only people authorized by the Statistics Act 1975 are allowed to see data about a particular person, household, business, or organization, and the results in this paper have been confidentialized to protect these groups from identification.

Careful consideration has been given to the privacy, security, and confidentiality issues associated with using administrative and survey data in the IDI. Further detail can be found in the privacy impact assessment for the Integrated Data Infrastructure available from www.stats.govt.nz.

The results are based in part on tax data supplied by Inland Revenue to Statistics NZ under the Tax Administration Act 1994. This tax data must be used only for statistical purposes, and no individual information may be published or disclosed in any other form, or provided to Inland Revenue for administrative or regulatory purposes.

Any person who has had access to the unit-record data has certified that they have been shown, have read, and have understood section 81 of the Tax Administration Act 1994, which relates to secrecy. Any discussion of data limitations or weaknesses is in the context of using the IDI for statistical purposes, and is not related to the ability of these data to support Inland Revenue's core operational requirements.

Macro and Micro Consumption Drivers in the Wake of the COVID-19 Pandemic*

Elisa Guglielminetti and Concetta Rondinelli
Bank of Italy

In this paper, we investigate the drivers of household consumption in the wake of the COVID-19 pandemic, focusing on Italy. We combine a macro and a micro approach and find that the deterioration in economic conditions can only explain about half of the slump in consumption, both at the aggregate and at the individual level; increased income and health risks significantly affect spending decisions, with heterogeneous impacts across expenditure categories and household types. At the micro level, the unemployed reduce spending for economic reasons, while the self-employed are mainly discouraged by health risks and uncertainty; restrictive measures play a minor role.


JEL Codes: D14, D15, E21.

1. Introduction

Following the outbreak of COVID-19, Italian private consumption fell dramatically (by about 10 percent), but the contraction in disposable income, which was sustained by the stimulus package implemented by the government since the outbreak of the pandemic, was somewhat smaller (about 3 percent). The saving rate then spiked and touched historically high levels, going above 20 percent in spring 2020. The spectacular increase in saving during the COVID-19 crisis was rather unusual, and seemed to contradict standard economic

*The views expressed herein do not reflect those of the Bank of Italy. We would like to thank, for their useful comments and suggestions, Francesco D'Acunto, Giuseppe Fiori, Luigi Guiso, Steffi Huber, Tullio Jappelli, Stefano Neri, Michaela Pagel, Daniele Terlizzese, and Michael Weber. Corresponding author: Concetta Rondinelli, concetta.rondinelli@bancaditalia.it. Bank of Italy, Economic Outlook and Monetary Policy Directorate, via Nazionale 91, Roma, 00184, Italy.

Figure 1. Changes in Household Consumption and Saving Rate between 2019 and 2020



Source: Federal Reserve Economic Data (FRED) database and Eurostat.
Note: Consumption is the percentage difference between the final consumption expenditure of households from the annual national accounts in 2019 and 2020. Saving is the change in the gross household saving rate from quarterly sectoral accounts between 2019 and 2020.

models that suggest that income falls more than consumption during recessions, as households attempt to smooth at least part of the negative shock, resulting in decreased saving rates (see Rodano and Rondinelli 2014 for the recent global financial crisis and sovereign debt crisis). Among the euro area countries, Italy, together with Spain, recorded the sharpest drop in consumption and the strongest increase in the saving rate (Figure 1); the unusual rise in saving affected the United States as well.

Given that the shock was exogenous, unexpected, and common to many countries, the recent health crisis is an ideal setting in which to study to what extent household spending is shaped by actual changes in economic fundamentals rather than by other, more subjective factors. During a health crisis the relevant household beliefs regarding consumption and saving choices include not only expectations about future income—most likely endogenous to the current economic

condition—but also other factors, such as the fear of infection and the desire to save more to face the considerable uncertainties about the evolution of the epidemic and its economic consequences. This makes pandemics different from standard recessions and likely to have more persistent macroeconomic effects (Jordà, Singh, and Taylor 2022), but it also provides useful variation in the motivations for cutting consumption that allows the role of subjective expectations for household spending to be identified.

In this paper, we investigate how economic conditions and other pandemic-related factors shaped consumption patterns in Italy during the COVID-19 pandemic, and we look into the contribution of these drivers across expenditure categories and household types. Several factors can be held responsible for these consumption and saving patterns. The first factor relates to the classical economic fundamentals: the fall in disposable income and the job losses may have induced households to cut spending; at the same time, the effect of economic hardship on savings is unclear, as liquidity constraints may have impaired the possibility of putting resources aside irrespective of the desire to save. Second, households might want to increase their saving buffer for precautionary reasons, because they were more uncertain about the evolution of their economic situation or because they perceived a higher health risk. Third, lockdown policies prevented some kinds of expenditure (restaurants and travel above all), generating forced savings. Fourth, the risk of infection may have discouraged households from consuming certain types of goods and services that require social contacts. All these factors likely played a role, albeit with different intensities for each spending category and for different types of households.

To tackle this issue, we combine a macro and a micro approach. On the macro side, we estimate a consumption equation à la Ando and Modigliani (1963) by regressing private consumption on its traditional determinants, such as income, wealth, interest rates, and expectations. We interpret the unexplained part of the consumption drop as being driven by pandemic-related factors not explicitly included in the regression, such as the fear of infection, lockdown policies, and increased uncertainty about the future. We find that standard economic factors can only explain about half of the slump in consumption, but the relative importance of each driver varies substantially across expenditure categories. As expected,

pandemic-related factors are more relevant for those items directly affected by lockdown policies and for which the perceived risk of infection is higher, such as restaurants, hotels, recreational activities, and clothing.

The macro approach allows us to distinguish the standard determinants of consumption relating to economic reasons, precautionary motives, and uncertainty, but it cannot single out the impact of the fear of infection and lockdown policies; by construction, these are included in the residual of the model. The pandemic event was a quite new type of recession in our sample, so time-series data do not provide us with sufficient variation to disentangle the impact of these specific factors. The macro framework therefore calls for an integration with micro data, which allows us to explore spending patterns by exploiting the cross-sectional variation. Hence we leverage data from the Bank of Italy's Special Survey of Italian Households (SSIH hereafter), which was launched after the onset of the COVID-19 pandemic. The survey contains questions on the respondent's financial situation, occupational status, saving possibilities, consumption choices, and expectations on future income changes and labor market prospects, which mimic the standard determinants of consumption used in the macro approach. Moreover, the survey asks questions about the changes in and motivations for spending behavior during the pandemic. Therefore, this unique data source allows us to build individual-level measures of the financial situation, uncertainty, and fear of infection, and to analyze their contribution to the expected evolution of consumption. Moreover, we capture the impact of containment measures by leveraging heterogeneous restrictions between Italian regions during the autumn of 2020.

To strengthen the link between the macro and micro approaches, the households interviewed in the SSIH were asked about spending on goods (non-durable, durable, and semi-durable) and services over the following three months. More precisely, the categories relating to food, clothing and footwear, home goods, and services in the SSIH were chosen to line up with those available at the aggregate level; this ensures comparability between the macro and micro approaches for the main outcome variable.

In line with the macro approach, the microeconomic evidence shows that the deterioration in economic conditions can only explain about half of the slump in consumption. Apart from economic

reasons, the fear of infection is the largest contributor to the probability of cutting spending in the following three months, followed by heightened uncertainty about the future due to economic and health concerns; once the previous factors have been controlled for, restrictive measures do not play a significant role. The micro approach can provide interesting insights into the importance of different pandemic drivers across households' characteristics, such as the occupational status and the expected possibility of saving: fears of infection and uncertainty account for a larger share of the drop in consumption for self-employed household heads; the fall in expenditure for unemployed household heads, instead, depends mostly on economic reasons. Moreover, uncertainty is a significant factor only for liquidity-constrained individuals, namely those who do not expect to be able to put resources aside in the following year.

The rest of the paper is organized as follows. Section 2 reviews the literature on consumption during the COVID-19 crisis and discusses the fiscal support Italian households received during the pandemic. In Section 3 we present the estimates of the aggregate consumption function by expenditure category. In Section 4 we describe the Bank of Italy's Special Survey of Italian Households and investigate the reasons for cutting spending using micro data. Section 5 concludes.

2. Consumption Patterns during the COVID-19 Pandemic

Our paper contributes to the long-standing research on consumption behavior, which acknowledges the important role of consumption risk and subjective expectations for explaining current spending decisions (Christelis et al. 2020b). In this paper, we exploit the unexpected and unprecedented nature of the COVID-19 crisis to quantify the importance of different factors in the response of consumption. Moreover, we can assess the relative importance of health and income risk by leveraging the individual-based measures constructed from the household survey.

We also provide novel evidence on consumption and saving patterns after the outbreak of COVID-19, which has been at the center stage of the academics' and policymakers' agenda since the burst of the pandemic. High-frequency bank and credit card transaction data show that households reacted to the spread of the virus with sizable

cuts in spending that were not evenly distributed across consumption categories (Andersen et al. 2020, Baker et al. 2020, Chetty et al. 2020, Cox et al. 2020). Furthermore, these cuts were much larger than those that could be explained by the job and income losses occurring at the same time and involved households across the whole income distribution (Chetty et al. 2020, Cox et al. 2020).¹ This evidence suggests that consumers' behavior during the pandemic was affected not only by economic factors but also by restrictive policies, infection concerns, and possibly by a precautionary attitude reinforced by the increase in both income and health risks. However, it is hard to discriminate between the latter factors: indeed, the extent to which spending on non-essential items was influenced by formal stay-at-home orders (lockdowns) is controversial (Andersen et al. 2020, Baker et al. 2020, Chetty et al. 2020, Alexander and Karger 2023). Using a different methodology applied to Italy, our work complements this evidence, highlighting the importance of the fear of contagion and precautionary motives for consumption choices, independently of government-mandated restrictions.

Finally, our work is closely connected to the literature that exploits household surveys carried out during the pandemic (Christelis et al. 2020a, Coibion, Gorodnichenko, and Weber 2020b, Coibion et al. 2021, Hodbod et al. 2021, Immordino et al. 2022). Based on the Consumer Expectations Survey (CES) conducted in the six largest euro area economies, Christelis et al. (2020a) find a strong effect of financial concerns due to COVID-19 on spending on non-durables, consistent with an important role of precautionary motives during the first peak of the pandemic and its immediate aftermath. Our results confirm and complement this evidence, as we highlight the importance of both economic reasons and idiosyncratic uncertainty for consumers' behavior. Our paper is close to that of Hodbod et al. (2021), who analyze the role of additional motivations for cutting consumption, such as the worries about the infection risk and permanent shifts in households' tastes. Like them,

¹As pointed out by Cox et al. (2020), the spending drop in March 2020 is roughly eight times larger than the average household credit card spending drop in the first month of unemployment for unemployment insurance recipients in normal times, as estimated by Ganong and Noel (2019) on U.S. data.

we quantify the importance of different drivers for consumption patterns, but we also provide evidence of significant heterogeneity across household categories. Furthermore, our macro approach provides an estimate of the impact of pandemic-related factors on aggregate consumption and offers complementary evidence on their heterogeneous importance for different expenditure categories.

Household spending was deeply affected by the pandemic, despite the large-scale fiscal measures implemented by the governments around the world, such as direct payments to individuals in the form of stimulus checks (Coibion, Gorodnichenko, and Weber 2020a) or temporary cuts in consumption taxes (Bachmann et al. 2021). Moreover, central banks adopted large-scale unconventional measures, including large-scale asset purchase programs (D'Acunto et al. 2021). In Italy, public transfers to households rose by more than 10 percent in 2020, contributing to limiting the fall in household disposable income, which only contracted by 2.7 percent (against a GDP drop of 9 percent); the decrease in income from self-employment was more marked (more than 10 percent).

The SSIH provides additional granular information on the evolution of household income and fiscal support measures. One-third of households interviewed in November 2020 asserted that their income declined in 2020 compared with the previous year; this share rises to over half among households headed by a self-employed worker or by someone unemployed (Rondinelli and Zanichelli 2021a). The same data suggest that between March 2020 and April 2021 more than 40 percent of Italian households had access to at least one form of income support for workers (wage integration, benefits for the unemployed and self-employed or freelancers, and other bonuses) or for households (Reddito di cittadinanza, Reddito di emergenza, and baby-sitting vouchers); of these, about one-third stated that they benefited from two or more measures.² Moreover, other policies such as the extension of wage integration schemes (Cassa integrazione guadagni, or CIG), the freeze on dismissals for economic reasons,

²In April 2019, a new minimum income scheme (Reddito di cittadinanza) was introduced in favor of households suffering economic hardship. This measure was temporarily expanded by the introduction of the emergency income (Reddito di emergenza) that was established by Decree Law 34/2020 and renewed by Decree Laws 104/2020, 137/2020, and 41/2021. Access to this measure was conditional on the possession of certain income and capital requirements.

and business support measures contributed to limiting the decline in the number of employed persons.

Not only was the evolution of income and fiscal support measures heterogeneous across the household distribution, but the spending response was as well. In 2020, Italian consumption decreased for all households, but more markedly for those whose head was an entrepreneur or a self-employed and for the most affluent ones (Istat 2021), as they generally spend a greater part of their budget on the consumption items hit the most by the pandemic.

In this paper, we take into account the government's support measures both in the macro and in the micro approach. At the aggregate level, fiscal transfers are included in household disposable income, which we consider in the macro regression as one of the main determinants of household consumption. At the micro level, we know whether households have received any kind of income support and we can control for it (see Section 4.3).

3. The Macroeconomic Approach

3.1 *The Econometric Strategy*

In this section we investigate the consumption pattern during COVID-19 in Italy from a macro perspective, by estimating a consumption equation à la Ando and Modigliani (1963), as adapted to the Italian economy by De Bonis et al. (2020). Our goal is to broadly quantify to what extent standard drivers can explain consumption dynamics in 2020, rather than estimate the precise elasticity of household spending to each factor by taking into account their long-term relationships. Hence we estimate a simple OLS regression³:

$$c_t = \alpha + X_t' \beta + \varepsilon_t, \quad (1)$$

where c_t is the log of household consumption in real terms, α is a constant, and X_t is a vector of covariates including the log of permanent income, income volatility, the log of real wealth, the log of financial

³De Bonis et al. (2020) instead adopt a vector error correction model, as they are interested in providing precise estimates of the elasticity of consumption to different components of wealth using a longer sample period.

wealth, the interest rate, hours worked, unemployment expectations, and uncertainty.⁴ In this way, we control for all the main determinants of consumption, which we can group in the following categories: (i) permanent income; (ii) current income, proxied by hours worked; (iii) wealth (both real and financial); (iv) the intertemporal rate of substitution, proxied by the interest rate; (v) unemployment expectations; and (vi) consumers' disagreement about the evolution of the economic situation.⁵ The evolution of permanent income is given by the four-quarters moving average of the household real disposable income; we also control for its volatility, computed as the standard deviation of income growth on a rolling backward window of two years.⁶ We capture current income through hours worked, which provides a better signal of the cyclical position of the economy compared to employment-based measures. This is of primary importance for 2020, when the government put restrictions on layoffs to safeguard jobs during the pandemic crisis.⁷ Real wealth is the value of dwellings owned by households and thus depends on the evolution of housing prices. Financial wealth is the collection of net financial assets held by households, both liquid and illiquid. We use the same definition of De Bonis et al. (2020): liquid assets include deposits, bonds, mutual funds, and quoted shares net of total liabilities, whereas illiquid financial wealth is the sum of unquoted shares and other equity plus holdings of insurance and pension fund instruments. We include the long-term interest rate at 10-years maturity to control for borrowing conditions. Finally, we introduce other two

⁴Appendix A contains a detailed description of the data used in this section.

⁵Along the lines identified by De Bonis et al. (2020), we tested alternative specifications distinguishing liquid from illiquid wealth and further splitting liquid wealth into assets and liabilities. In other robustness checks we included a credit conditions index and the log of the interest rate, thus allowing for potential non-linear effects of this variable. None of these regressors had a significant effect on total consumption, and we thus excluded them from the baseline specification. Results are available upon request.

⁶Using different definitions of permanent income, such as the 8- or 12-terms moving average of household disposable income, leads to a reduced significance of the elasticity of consumption to this factor and to an increase in the magnitude of the coefficient on hours worked. The other results are almost unchanged.

⁷Even in normal times the fit of the consumption equation improves moving from an employment-based measure to the one based on hours worked.

variables that may hold back spending, namely consumers' expectations and their dispersion. Unemployment expectations, taken from the consumer survey of the European Commission, account for consumers' mood about the overall evolution of the labor market. A rise in unemployment expectations, everything else equal, suggests rising concerns about the economic situation that should induce households to save more. Lastly, we build a measure that reflects the dispersion of consumers' responses in the survey of the European Commission. We consider the principal component of such indices based on the expectations about their own financial situation, the national economic situation, and the labor market.⁸

Regarding the dependent variable, we consider several alternatives. First, we estimate Equation (1) on real consumption of resident households, from national accounts data. Then we consider a breakdown by six expenditure categories: (i) hotels and restaurants, (ii) recreation and culture, (iii) clothing and footwear, (iv) furnishing and house equipment; (v) personal care; (vi) food and beverage.⁹ This distinction can provide interesting insights about the heterogeneous impact of the pandemic and the importance of different motives underlying spending patterns. Government restrictions should have the strongest impact on the travel and accommodation industry. Health concerns may possibly regard a wider range of goods and services that require social interactions. Changes in households' habits—above all, the surge in teleworking—may have also shifted their preferences towards different goods, like furnishing and house equipment. As shown by other works on U.S. data (Chenarides et al. 2021), food and beverage consumption could be positively affected by substitution effects between restaurants and home production and, at the outbreak of the pandemic, by consumers' panic about possible disruptions in the food supply chain. Moreover, food consumption should be the least affected by the plausible intensification of precautionary attitudes.

⁸The dispersion index for question q is computed as $I^q = \sum_{j=1}^{j=5} F_j^q (1 - F_j)$, where F_j^q is the cumulated frequency of responses of type j at question q , where j ranges from "very negative" to "very positive." Hence, the index ranges from 0 (no dispersion) to 1 (maximum dispersion). The indices are also standardized to have 0 mean and unitary standard deviation.

⁹For personal care spending we take the data from Confcommercio.

3.2 *The Results*

We estimate the model by OLS on quarterly data over the period 2001:Q2–2019:Q4 (74 observations).¹⁰ We intentionally exclude the pandemic period from the estimation sample because it likely represents a structural break in the historical correlations among the selected variables that would distort the estimated coefficients if not properly taken into account. Our aim is to show how much of the drop in consumption can be explained by standard drivers according to their pre-pandemic relationships.¹¹ In this way the unexplained part of the spending cut can be interpreted as related to the pandemic, either directly through the fear of contagion, government-mandated restrictions, and heightened uncertainty about the future due to economic and health concerns, or indirectly through an amplification of the importance of standard drivers and changes in households' preferences. The estimates show that the income channel is stronger for non-essential items such as restaurants and furnishings, while food consumption is less sensitive to changes in permanent income or hours worked.¹² The impact of real and financial wealth is overall modest. The interest rate is negative and significant for total consumption, restaurants and hotels, and personal care. Unemployment expectations and consumers' disagreement about economic prospects weight negatively on all categories except furnishings and food.¹³

¹⁰The choice of the sample period is motivated by data availability. Household real disposable income is available since 1999:Q1, hence the quarter-on-quarter (q-o-q) growth rate starts in 1999:Q2 and the first observation of the volatility of income is two years later.

¹¹Although the model is deliberately simple, it captures consumption dynamics considerably well in the pre-pandemic period. To see this, we estimated the model with data up to 2015 and obtained out-of-sample forecasts for the period 2016–20. The root mean squared forecast error (RMSFE) of our model on the year-on-year (y-o-y) growth rate of total consumption is 0.4, considerably smaller than the unconditional standard deviation of the time series (1.3 percent) and the RMSFE of a simple time-series model (an autoregressive process with two lags, selected to maximize information criteria), which yields an RMSFE of 0.7.

¹²The estimates of the model coefficients are reported in Appendix B, Table B.1.

¹³Food consumption is positively affected by unemployment expectations and negatively affected by consumers' disagreement.

To assess the contribution of the different factors to consumption dynamics during the pandemic, we apply the estimated β coefficients to the y-o-y changes of the covariates in the four quarters of 2020. In this way, we get the contribution of the different drivers to the predicted y-o-y change in household consumption.¹⁴ We expect a modest impact of permanent income and wealth, as they are slow-moving variables, as well as of financial conditions, which could not be eased much further given the low interest rate environment in place already before the outbreak of the pandemic. In contrast, changes in hours worked and in the survey-based measures, which are more responsive to cyclical conditions, could have a stronger influence on spending patterns in 2020.


As expected, given the sharp economic contraction caused by the spread of the virus, the income effect proxied by hours worked explains about 70 percent of the predicted drop in spending in 2020 (Figure 2). Consumers' worries and divergent expectations about economic and labor market prospects also play a relevant role, accounting for more than 20 percent of the consumption cut. The accommodative monetary policy stance, which maintained interest rates low (here included in the "Other" factors), gave instead a small positive contribution. Notice that the model can only explain half of the dramatic slump in consumption in the first half of 2020 and about 60 percent of the y-o-y variation in 2020:Q3, while predicted consumption growth is much closer to the realized figure for 2020:Q4.¹⁵

The unexplained share of the drop in spending over the year is due to the peculiar features of the COVID-19 shock, which cannot be captured by the standard determinants of consumption included in the model. As in Dossche and Zlatanos (2020), we can interpret the residual of the estimated model as the contribution of pandemic-related factors, such as fears of contagion, restrictive

¹⁴We report the y-o-y instead of q-o-q changes to make more transparent the drop in spending compared to the previous year; in fact, for some consumption categories the q-o-q change is milder in Q2 than in Q1 just for a level effect.

¹⁵The difference between realized and predicted consumption growth rates (in absolute values) amounts to 4.1 percentage points in Q1, 8.4 points in Q2, 2.8 points in Q3, and 1.9 points in Q4. Realized and predicted growth rates for all consumption categories and for all quarters are reported in Appendix B, Table B.2.

Figure 2. Household Consumption Growth and Main Drivers




Source: Our computations on Istat, Bank of Italy, and ECB data.
Note: The solid line is the y-o-y realized growth in consumption. The dashed line represents y-o-y consumption growth predicted by regression (1). The bars represent the contribution of different factors to the predicted y-o-y consumption growth, expressed in percentage points. The category “Other” includes the volatility of income and the interest rate.

policy measures, and heightened uncertainty about the future associated with either economic or health concerns. Although this aggregate approach does not allow us to disentangle these factors, we can gather additional information by observing how the share of the unexplained spending drop varies by expenditure category (Figure 3).¹⁶

Given the nature of the shock, we expect the difference between predicted and actual growth rates in total consumption to be higher for the consumption categories which suffered the most from COVID-19. Focusing on the spring months of 2020, the residual is indeed very sizable for restaurants and hotels and for recreational

¹⁶Figure B.1 in Appendix B represents the unexplained drop in spending in all quarters.

Figure 3. Unexplained Consumption Growth in Macro Regressions (realized – predicted)



Source: Our computations on data from Istat, Bank of Italy, ECB, and Confindustria.

Note: For each consumption item we estimate the regression (1) and compute the difference between realized and predicted y-o-y consumption growth (residual).

and cultural activities, whose expenditure decreased by almost 70 and 40 percent, respectively, compared to 2019:Q2 (in 2020, as a whole, spending in these categories declined by 40.6 and 22.5 percent, respectively). Additionally, the standard drivers cannot explain one-third of the drop in clothing and footwear in 2020:Q2 (about half in the first three quarters of the years). Expenditure in furnishings and house equipment reduced by about 7 percent in 2020, much less than what one would have expected on the basis of the historical correlations with hours worked and permanent income; the residual is hence positive for this category. This result is probably due to the change in life style determined by the COVID-19 shock, which forced households to spend more time at home, inducing them to reallocate some of their expenditures to furnishings. Spending in personal care items fell by 6.3 percent; in this case, however, the model would have predicted a milder drop, as evident from the negative residuals. Lastly, food and beverages consumption is the only spending category which registered a positive variation in 2020 (1.9 percent),

higher than the model prediction.¹⁷ In conclusion, the model overestimates spending in restaurants and hotels, recreation and culture, clothing, and personal care; at the same time, it underestimates the expenditure in furnishings and food consumption. Notice that the “abnormal” fall in consumption was larger but not limited to the first half of the year, characterized by a strong lockdown, and to restaurants and hotels, the category mostly affected by the policy measures. This suggests that other factors, like the fear of contagion and uncertainty about the future, do play a role in the pattern of aggregate consumption. In what follows we dig deeper into these factors using micro data.

4. The Microeconomic Approach

Complementing the analysis with micro data allows us to overcome two limitations of the macro approach. First, we can dig deeper into the share of the exceptional drop in consumption not explained by the standard determinants and disentangle the relative importance of pandemic-related factors. Secondly, by building individual-level measures of such drivers, we can study how their contribution varies with households’ characteristics.

4.1 *The Bank of Italy’s Special Survey on Households*

Since April 2020, soon after the outbreak of the COVID-19 pandemic, the Bank of Italy has been conducting a Special Survey of Italian Households to collect information on the impact of the epidemic on the financial situation and expectations of Italian households.¹⁸ Apart from the main demographic variables (sex, age, education, area of residence), the survey contains questions on the respondent’s financial situation, occupational status, saving possibilities, consumption choices, and expectations on future income changes and labor market prospects.

¹⁷The fit of the macro regression is very high, with an adjusted R-squared above 0.9 for all categories except clothing and footwear (0.66). This is likely due to the higher volatility of the latter category (2.9 in the pre-pandemic period, against 1.3 for total consumption), which makes it harder to predict.

¹⁸See Appendix C for details on the SSIH.

In this work we use mainly the third wave of the survey that was conducted in the autumn of 2020, when the restrictive measures for the containment of the second pandemic wave were becoming more and more stringent, though less restrictive compared to April. Survey questions elicit consumption and saving expectations. In particular, households are asked about their saving intentions, i.e., whether over the next 12 months they plan to spend less than the entire yearly income and succeed in saving (see Appendix C for the exact wording). More than 40 percent of households expect to spend less than their annual income in the next 12 months (Table C.3). This share rises to 67 percent among households that expect their income to increase, but it is non-negligible (just under 20 percent) also among households that expect their income to fall. Two considerations are in order. First, this question provides information on the extensive margin of saving, allowing us to distinguish (expected) savers from non-savers. However, it does not tell anything about the intensive margin of saving, namely if the individual expects to save less or more compared to the past (conditional on being a saver). Second, the saving response is the net effect of income and consumption developments. Therefore, respondents may expect to succeed in saving for very different reasons: because they rely on a safe income which more than satisfies their usual needs, because they actively decide to save some resources to face unexpected events, or rather because they choose or are forced to cut consumption—for instance, because they fear being infected or because policy restrictions reduce the availability of some goods and services. Hence, to shed light on the channels through which the pandemic altered the saving behavior of Italian households, we start by focusing on the changes in spending decisions, which are more directly affected by the spread of the virus and fully under the households' control.

We thus exploit the question about expected spending in goods (non-durable, durable, and semi-durable) and services over the following three months.¹⁹ A little less than one-third of Italian households think that they will reduce their consumption; about half of them plan to cut their spending by more than 20 percent (Table C.3). The decrease is seen largely among households in the regions most

¹⁹See Appendix C for the exact wording.

exposed to the health emergency. The reasons for cutting expenditure seem to be not only economic: over one-fourth of individuals that expect their income to increase in 2021 plan to cut spending.

Additional evidence on the reasons for cutting expenditure comes from a specific question introduced in the third wave of the survey, where households compare their spending in the previous month to that before the COVID-19 pandemic.²⁰ The question is asked separately for different categories: (i) food, (ii) clothing and footwear, (iii) personal services, (iv) furniture, and (iv) bars, restaurants, and hotels. Consistent with the macroeconomic evidence, around 80 percent of households report that they have patronized some establishments less regularly—“hotels, coffee bars, and restaurants”—and have made less frequent purchases in clothing stores than prior to the pandemic; about two-thirds have cut spending on beauty and personal care services.²¹ The decline in these categories of consumption items is greater in the regions hit the most by the public health emergency; starting from the beginning of November 2020, Italian regions were differentiated on a weekly basis according to the degree of restrictions in place, identified either by a yellow, orange, or red color (with increasing severity).²² If households report to have reduced spending in at least one category of clothing and footwear, personal services, bars, hotels, and restaurants, they are asked the

²⁰Question 1 of the COVID-19 module (Appendix C).

²¹In the case of services, it is more realistic to associate less frequent visits with lower spending. In relation to the categories of food and clothing, instead, less frequent visits to the physical stores may be accompanied by larger volumes of purchases (go less, buy more) or by the use of other shopping modes that could compensate for less in-store shopping, such as curbside pickup, online shopping, and shopping by others (Ardizzi, Nobili, and Rocco 2020).

²²Our data suggest that the decline in consumption was sharper in red and orange regions, which featured both high infection rates and more stringent restrictions, compared to yellow regions. There is also external evidence that supports this claim. The drop in overall and service consumption in 2020 at the regional level is negatively correlated with the cumulative sum of infections (rescaled by the population) in the same region; results holds when considering hospitalizations or deaths instead of infections. Moreover, Emiliozzi, Rondinelli, and Villa (2023), using high-frequency credit card data, show that (i) overall consumer spending dropped by more in the regions most hit by the pandemic and (ii) consumer spending dropped by more in those categories more exposed to the risk of infection.

motivations of this choice.²³ Just under half of the households say that the reduction in consumption is due to diminished economic resources (Table C.3). Among the other reasons given, fear of infection is the main driver, regardless of the severity of the restrictions imposed in the interviewees' regions of residence. The forced reduction in spending due to measures taken to contain the epidemic has a lesser impact overall, although it is greater for households that reside in "red zones." The average score assigned to putting money aside for unexpected events is 29, with a higher importance in "yellow zones."

4.2 Reasons for Cutting Expenditure

Despite being primarily a qualitative survey, SSIH allows us to estimate a micro version of Equation (1), with the goal of disentangling the drivers of the expected consumption pattern over the following three months at the individual level. We take advantage of the rich set of information included in the survey to control for both standard and pandemic-related determinants of consumption choices, all considered at the individual level. Among the standard factors, we control for (i) age and education, which capture the stage of the individual's life-cycle; (ii) actual income, proxied by the occupational status; (iii) liquid wealth, captured by the difficulties in making ends meet;²⁴ and (iv) expectations about future resources, which are assumed to be a function of individual expectations on income and labor market developments.

Moreover, we are interested in disentangling the other channels behind the exceptional drop in consumption and increase in saving. To do so, we build individual-level measures, which are either exacerbated by the pandemic (like economic reasons and uncertainty about the future) or directly linked to it (like the fear of contagion and containment measures). The three questions on the reasons that induced households to patronize some establishments less regularly help in this regard. Rather than using directly these responses as

²³Questions 2–3 of the COVID-19 module (Appendix C).

²⁴Households reporting to have difficulties in making ends meet are three times more likely to have not enough financial resources to face essential expenses for three months compared to households who easily make ends meet.

regressors, which would limit the sample size, we first try to isolate which other variables are mostly associated with the four reasons.

Households may want to put money aside for unexpected events to smooth the uncertainty they will face in the future. This uncertainty might be primarily affected by labor market events. For this reason we construct a measure of individual uncertainty as a dummy variable equal to one if the self-reported probability of (i) losing the job is higher than 25 percent for employed household members and (ii) finding a job is lower than 25 percent for unemployed individuals. Notice that, by construction, the uncertainty variable is only available for households who have at least one member in the labor force and hence it is missing for almost 80 percent of households whose head is a pensioner.²⁵ In what follows we will conduct robustness checks to take into account the possible selection bias generated by our definition of uncertainty.

Restrictive measures are more stringent in the regions most severely hit by the epidemic at the time of interview, the so-called red and orange zones.²⁶

Households that decided to cut drastically spending in all types of non-essential goods and services requiring social interactions—like hotels, bars, and restaurants—and purchases in physical stores selling clothing and footwear or furniture and appliances, might be signaling a strong fear of contracting the virus. For this reason we create a variable indicating fear as a dummy variable equal to one when the household answers “1” to questions 1.2, 1.3, and 1.5 in Appendix C.1 (FEAR). As a robustness check, we also construct an alternative measure using a different survey question: in this case we consider how spending would vary depending on the number of infections per day in the own region. Fearful households would eliminate or strongly reduce their spending in hotels, coffee bars, and restaurants even in a more favorable epidemiological scenario. Our alternative measure of fear is thus a dummy variable equal to one if the

²⁵In terms of observable characteristics, the pensioners for whom it is possible to compute the uncertainty variable are significantly more males, married, and with a higher number of household members, who thus have a higher probability of being in the labor force. There are no significant differences in terms of age or education.

²⁶See Section C.4 in Appendix C for the classification of Italian regions in November 2020.

household decides to eliminate or strongly reduce spending in hotels, coffee bars, and restaurants even when in the region of residence there are fewer than 10 new cases per day (FEAR CASES).²⁷

To select the households' characteristics associated the most with the reasons for cutting expenditure, we estimate probit regressions for the probability of reporting any of these motivations as particularly relevant. For the probit model related to economic reasons, the dependent variable is a dummy equal to one when the household reported to have reduced purchases because of lower financial resources (question 2 of the COVID-19 module). For the other motivations, we instead created dummy variables which take a value of 1 if the score that the respondent assigns to a given motivation is higher than the median.²⁸

Table 1 reports the results of the probit models, where each of the four reasons for cutting expenditure is regressed against age, occupational status, making ends meets, expectations on income and labor market, and the three variables described above that we expect to be correlated with the motivations other than the economic ones. First, cutting expenditure because of economic reasons appears, as expected, to be positively correlated with occupational status: self-employed are about 22 percent more likely to cite economic reasons compared to pensioners, in line with the evidence that the pandemic affected self-employment the most. Those experiencing difficulty in making ends meets cite more frequently economic reasons compared with those not experiencing that difficulty. Additionally, compared to those with a more optimistic view, households with less favorable expectations for the labor market and income mention economic reasons 10 percent more often. Economic reasons are also positively correlated with the fear of contagion and uncertainty. Second, in the model for the probability of mentioning the fear of contagion as an important reason for reducing consumption, the variable FEAR is highly significant, as expected; the same evidence applies with a narrower definition of fear (FEAR CASES);

²⁷See question 4 in Section C.1 of Appendix C. The question is the same irrespective of the size of the region of residence, but 10 new cases is considered low even in small regions.

²⁸In theory we could classify individuals depending on the prevalent motivation; however, cases where the respondent assigns equal weights to at least two motivations are not infrequent.

Table 1. Determinants of the Reasons for Cutting Expenditure (probit models)

	Economic Reasons (1)	Fears of Contagion (2)	Fears of Contagion (3)	Restrictive Measure (4)	Deal with Unexpected Events (5)
Age	-0.001 (0.011)	-0.012 (0.014)	-0.012 (0.014)	0.008 (0.013)	-0.000 (0.013)
Age Sq.	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)
Employee	0.108 (0.075)	-0.182** (0.085)	-0.206** (0.084)	-0.004 (0.092)	0.046 (0.088)
Self-Employed	0.223*** (0.076)	-0.129 (0.091)	-0.144 (0.090)	-0.017 (0.102)	0.060 (0.100)
Unemployed	0.130 (0.088)	-0.145 (0.110)	-0.161 (0.107)	-0.040 (0.122)	-0.051 (0.120)
Making Ends Meet (Great Difficulty)	0.494*** (0.041)	0.012 (0.093)	0.002 (0.092)	-0.185** (0.089)	0.036 (0.091)
Making Ends Meet (Difficulty)	0.391*** (0.047)	-0.056 (0.061)	-0.052 (0.062)	-0.029 (0.062)	0.169*** (0.059)
Making Ends Meet (Easily)	0.098* (0.055)	-0.081 (0.054)	-0.078 (0.054)	-0.060 (0.055)	0.102* (0.054)
Expectation on Y (Decrease)	0.119*** (0.045)	0.006 (0.068)	0.005 (0.068)	-0.052 (0.063)	-0.009 (0.068)
Expectation on LM (Decrease)	0.093** (0.038)	0.021 (0.044)	0.028 (0.044)	-0.037 (0.044)	-0.037 (0.045)

(continued)

Table 1. (Continued)

	Economic Reasons	Fears of Contagion	Fears of Contagion	Restrictive Measure	Deal with Unexpected Events
	(1)	(2)	(3)	(4)	(5)
Red Zone	0.031 (0.048)	0.027 (0.055)	0.046 (0.054)	0.160*** (0.053)	-0.011 (0.055)
Orange Zone	0.072 (0.047)	0.071 (0.057)	0.066 (0.057)	0.077 (0.054)	0.061 (0.058)
FEAR	0.098** (0.041)	0.147*** (0.052)		-0.143*** (0.052)	0.050 (0.052)
Uncertainty	0.144*** (0.042)	0.035 (0.062)	0.040 (0.062)	-0.008 (0.062)	0.114* (0.061)
FEARS CASES			0.080* (0.044)		
Obs.	1,248	685	685	685	685
Pseudo R^2	0.194	0.037	0.031	0.035	0.022
Average Probability	0.48	0.46	0.46	0.59	0.54
Note: Marginal effects. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Our calculations from wave 3 of SSIH. Dependent variable is a dummy equal to one when the household reported to have reduced purchases because of lower financial resources (question 2 in Appendix C.1; column 1). For the other motivations, the dependent variable is a dummy variable which takes value of 1 if the score that the respondent assigns to a given motivation is higher than the median (question 3 of the COVID-19 module in Appendix C; columns 2-5). Reference category for education is college and beyond. LM indicates labor market, Y indicates income; reference category for expectation on Y and LM is increase or stay the same. Reference category for occupational status is pensioner; the unemployed category includes students, housewives, and others. Reference category for making ends meet is very easily.					

employees mention less often the fear of contagion compared to pensioners. Third, households mentioning that the drop in spending was related to the restrictive measures are more likely to live in red zones at the moment of the interview. We may be worried that our approach is not able to distinguish between households caring about the possibility of infection and those that are forced to reduce purchases because of the containment measures; however, the coefficient of FEAR is negative and significant, mitigating the concerns of a positive correlation between these two factors. Fourth, putting money aside for unexpected events is positively correlated with our uncertainty measure and the difficulty in making ends meet.

Thus, we can conclude that among the four reasons for cutting consumption, making ends meet, occupational status, and expectations are mainly correlated with the economic reasons; FEAR correctly captures the fear of contagion; the red zone is a good predictor for mentioning restrictive measures, and putting money aside to face unexpected events depends on the uncertainty on the own employment prospects, although the latter variable is also correlated with economic reasons.²⁹

4.3 *Disentangling the Drivers of Consumption*

Having described the survey variables that capture both the standard and the pandemic-related factors that could affect consumption choices, we can now estimate a micro version of the consumption function³⁰:

$$\begin{aligned} ncexp_i = & \beta' Z_i + \gamma' Y_i + \delta' W_i + \eta' EXP_i \\ & + \Theta' reason_i + \varepsilon_i, \end{aligned} \quad (2)$$

where $ncexp_i$ is a dummy variable equal to one for households expecting a decrease in their consumption for food, clothing and footwear, and home goods and services in the following three months (see Section C.4 in Appendix C) and Z_i is a deterministic component

²⁹Results are broadly unchanged when we estimate a multinomial logit with all the reasons (Table D.1 in Appendix D).

³⁰Other papers estimate a micro version of the consumption function for Italy (see, among others, Guiso, Paiella, and Visco 2006, Paiella 2007, and Paiella and Pistaferri 2017).

including both education and age, which captures different stages in the life-cycle consumption profile. Individual actual income Y_i is proxied by the occupational status (employee, self-employed, unemployed, pensioner), while liquid wealth (W_i) is approximated by self-reported difficulty in making ends meets. The individual expectations about income and the evolution of the labor market are other important determinants of spending decisions (EXP_i). Among the four reasons for cutting consumption, the economic ones are captured by income, liquid wealth, and the self-reported expectations about income and labor market developments. The other three reasons for cutting consumption ($reason_i$) are based on the variables identified in Section 4.2, namely the fear of contagion (FEAR), the restrictive measures (red and orange zone), and uncertainty.

In Table 2 we report the estimates of Equation (2), starting from a simple specification of a linear probability model and gradually enriching it.³¹ In column 1 we see that expenditure has an inverse U-shaped profile over the life cycle, as demonstrated by the coefficients in age and age squared. Those with none or primary or middle school education are more likely to cut consumption compared to the highly educated. Unemployed and self-employed are more likely to expect a fall in consumption compared to pensioners. In column 2 we insert a variable aiming at capturing household liquid wealth, i.e., the capability in making ends meets: those finding it very difficult are more likely to expect a drop in consumption in the following three months compared to more affluent households; the introduction of this variable makes the one related to education not significant. In column 3 we also include individuals' expectations on income and the labor market: as expected, those foreseeing a drop in income or a worsening of the labor market mention more frequently a spending cut. Because education is not significant once controlling for the difficulty in making ends meets and expectations, we drop it from the other specifications, which include the other potential reasons for cutting expenditure. The fear of contagion is positively related to the drop in spending, both using the narrow and the broad definition of fear of contagion (columns 5 and 6);

³¹We estimate Equation (2) using a linear probability model to facilitate the disentangling of the drivers of consumption. Results, available upon requests, are confirmed using a probit model.

Table 2. Assessing the Channels in Cutting Expenditure (linear models)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Age	0.013** (0.005)	0.011** (0.005)	0.014*** (0.005)	0.013*** (0.005)	0.011 (0.007)	0.010 (0.006)	0.010 (0.006)	0.010 (0.006)
Age Sq.	-0.000** (0.000)	-0.000** (0.000)	-0.000*** (0.000)	-0.000** (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
None or Primary	0.146*** (0.039)	0.075* (0.040)	0.056 (0.038)	0.056 (0.038)				
Middle School	0.075** (0.033)	0.027 (0.033)	0.016 (0.031)	0.016 (0.031)				
High School	0.052 (0.033)	0.027 (0.033)	0.016 (0.031)	0.016 (0.031)				
Employee	0.053 (0.035)	0.049 (0.035)	0.045 (0.033)	0.045 (0.033)	0.074 (0.057)	0.080 (0.055)	0.082 (0.055)	0.081 (0.055)
Self-Employed	0.124*** (0.043)	0.116*** (0.043)	0.059 (0.040)	0.059 (0.040)	0.084 (0.059)	0.091 (0.058)	0.091 (0.058)	0.090 (0.058)
Unemployed	0.156*** (0.038)	0.130*** (0.037)	0.103*** (0.035)	0.108*** (0.035)	0.092 (0.063)	0.074 (0.061)	0.076 (0.061)	0.075 (0.062)
Making Ends Meet (Great Difficulty)		0.206*** (0.036)	0.129*** (0.034)	0.136*** (0.033)	0.083*** (0.040)	0.071* (0.039)	0.073* (0.040)	0.073* (0.040)
Making Ends Meet (Difficulty)		0.198*** (0.031)	0.153*** (0.030)	0.161*** (0.029)	0.103*** (0.036)	0.100*** (0.035)	0.101*** (0.035)	0.101*** (0.035)
Making Ends Meet (Easily)		0.071** (0.031)	0.036 (0.030)	0.038 (0.030)	0.014 (0.036)	0.009 (0.035)	0.010 (0.035)	0.011 (0.035)

(continued)

according to the broader classification (FEAR), a household that attributes high importance to the fear of contagion is 23 percent more likely to reduce purchases. Individuals who are highly uncertain about their labor market prospects are significantly more likely to compress expenditure. The restrictive measures, i.e., being either in the red or orange zone at moment of the interview, are positively related to the fall in expenditure, but the coefficient is not statistically significant. Column 7 is our preferred specification of Equation (2), where we include standard demographic controls and reasonable proxies for the different motivations for cutting expenditure.

The question on spending in the previous month relative to before the COVID-19 pandemic, used to construct the reasons for cutting consumption, faces the issue that consumption patterns at the categorical level are seasonal and the month before the survey might be special. To control for these possible biases we conduct the analysis on different waves of the survey led in alternative months of the year (see Appendix C for a description of the survey period). Our results are unchanged when the analysis is conducted on wave 2 of the survey, on the pooled sample of wave 2 and 3, and on wave 4 (Table D.2 in Appendix D). Our results are also robust to the inclusion of other variables: the interaction of education with fear of contagion and uncertainty and the sector of activity turn out to be not statistically significant and leave the results unchanged. As discussed in the previous section, since the uncertainty variable presents some missing values, its inclusion reduces the sample size and may create a selection bias. By excluding uncertainty from our preferred specification, we find that several coefficients gain significance, including those related to non-economic variables (e.g., age, red and orange zones). This suggests that our uncertainty measure is also related to a broad category of unexpected events, beyond the economic ones.³²

As discussed in Section 2, more than 40 percent of households had access to at least one form of income support for workers or for households (Bank of Italy 2021c). The introduction of a dummy variable equal to one for households having received fiscal support to

³²Results on the robustness exercises are available upon request.

income by the Italian government in our preferred specification is not statistically significant and leaves our results unchanged (column 8 of Table 2).


Using our preferred specification of Equation (2) (column 7 of Table 2), we compute the contributions of the different reasons for cutting expenditure for the average household in the sample.³³ The results are represented by the bar “Total” in Figure 4. The predicted probability of an expected drop in spending is equal to 0.35, which compares with an unconditional probability of 0.33. Among the four motivations, we find that when controlling for all the main determinants of consumption in Equation (2), the contribution of economic reasons is more than 50 percent; the fear of contagion accounts for about 20 percent and the weight assigned to unexpected events and to restrictive measures is slightly higher than 10 percent, though we know from Table 2 that the coefficient on restrictive measures is not statistically significant.

4.4 *Heterogeneous Effects of the Pandemic*

To investigate the heterogeneous impact of the pandemic across households, it is interesting to perform the same decomposition of Section 4.3 over different categories. One of the most relevant dimension of heterogeneity is represented by the job status, as already shown by the descriptive statistics: overall, one-third of households assert that they have suffered a decline in income in 2020, but this percentage rises to over half when the household is headed by a self-employed worker or by an unemployed individual. Hence we separately estimate our preferred specification of Equation (2) on individuals holding different occupations (Table D.3) and report the contributions of the reasons for cutting expenditure for the average household in each sub-sample (Figure 4). We find substantial heterogeneity across households depending on their job status. Forty percent of households headed by a self-employed individual expect a

³³For instance, to compute the contribution of the fear of contagion, we multiply the coefficient for FEAR in column 7 of Table 2 (0.23) for the average value assumed by the variable FEAR in the estimation sample (0.26). We thus find a value of 0.06, which we compare to the sum of the contributions provided by economic reasons, fear of infection, restrictive measures, and precautionary motives, obtaining a relative contribution of FEAR of 0.21.

Figure 4. Contribution of the Reasons for Cutting Expenditure by Occupation



Note: Our calculations from wave 3 of SSIH. Contributions are computed at the sample means of each occupational category.

drop in expenditure: about half of the predicted fall is explained by fears of contagion, and more than one-third is explained by uncertainty on unexpected events, in line with the huge increase in volatility of small firms' stock returns during a recession (Perez-Quiros and Timmermann 2000). About half of unemployed households heads expect a decline in expenditure, mentioning economic reasons above all. The contributions of the reasons for cutting expenditure for employees are broadly in line with total households. For pensioners, precautionary reasons exert a negative effect on the probability of cutting consumption, which is reasonable given that they need to smooth these events on a relatively shorter period of time. Restrictive measures are mildly significant for the employees, the unemployed, and the pensioners.³⁴ As argued before, however, the regression for pensioners may suffer from a selection bias induced

³⁴Results are confirmed when instead of estimating a different equation for each occupational status we estimate a unique equation interacting each job status with the fear of contagion and with uncertainty.

by the inclusion of the uncertainty variable, which is only available for a subset of them. To address this issue we re-estimate the model by excluding the uncertainty variable. Results confirm that the predicted probability of a drop in expenditure is lower for pensioners than for other occupational categories and economic reasons are the largest contributors; however, the fear of infection and the restrictive measures also gain relevance (last bar of Figure 4).

When we adopt the same decomposition based on the age of the household's head, we find that the probability of cutting consumption is relatively higher for people aged 55–64 and lower for older respondents (more than 65 years old).³⁵ The contribution of restrictive measures is higher for the elderly; when excluding the uncertainty variable, which considerably reduces the sample size for this category, economic reasons become more relevant, as in the rest of the population. Restrictive measures provide a negative contribution for younger individuals: this could reflect the desire of the youth to resume spending as soon as tight restrictions are eased.

As already explained in Section 4.1, it is not straightforward to estimate Equation (2) using as a dependent variable the probability of being able to save resources over the following year.³⁶ However, to connect the drop in consumption with the saving expectations, we estimate our baseline regression (column 7 in Table 2), splitting the sample between those who plan to spend less than their yearly income and succeed in saving ($S = 1$) and those who do not (see Table 3). Splitting the sample into the two subgroups reveals that uncertainty and labor market expectations play a role in explaining consumption drop only for households not expecting to save, presumably reflecting their liquidity constraints: 40 percent have indeed some difficulty in making ends meet.³⁷

³⁵See Figure D.1 in Appendix D.

³⁶Only 10 percent of households expecting a fall in expenditure in the following three months foresee that they will be able to save some money in the following year.

³⁷Results are confirmed when considering only households in the panel component of SSIH who respond, in wave 4 (February–March 2021), to the question as to whether they accumulated savings in 2020 (Table D.5).

**Table 3. Connecting the Drop in Consumption
with Saving Increase (linear models)**

	Baseline	Saving = 1	Saving = 0
	(1)	(2)	(3)
Age	0.010 (0.006)	0.002 (0.009)	0.011 (0.009)
Age Sq.	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Employee	0.082 (0.055)	0.123 (0.081)	0.048 (0.075)
Self-Employed	0.091 (0.058)	0.076 (0.086)	0.077 (0.078)
Unemployed	0.076 (0.061)	0.169* (0.094)	0.023 (0.082)
Making Ends Meet (Great Difficulty)	0.073* (0.040)	0.115* (0.059)	0.042 (0.067)
Making Ends Meet (Difficulty)	0.101*** (0.035)	0.098** (0.046)	0.089 (0.064)
Making Ends Meet (Easily)	0.010 (0.035)	-0.044 (0.041)	0.059 (0.067)
Expectation on Y (Decrease)	0.286*** (0.029)	0.284*** (0.058)	0.253*** (0.035)
Expectation on LM (Decrease)	0.064*** (0.024)	0.035 (0.033)	0.075** (0.034)
Red Zone	0.044 (0.030)	0.058 (0.041)	0.028 (0.042)
Orange Zone	0.044 (0.030)	0.053 (0.043)	0.018 (0.042)
FEAR	0.228*** (0.027)	0.159*** (0.040)	0.272*** (0.037)
Uncertainty	0.115*** (0.028)	-0.017 (0.043)	0.172*** (0.036)
Constant	-0.312* (0.163)	-0.134 (0.229)	-0.295 (0.238)
Obs.	1,408	617	791
R ²	0.205	0.120	0.223

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Our calculations from wave 3 of SSIH. Dependent variable equal to one for households expecting a drop in expenditure in the next three months; in columns 2 and 3 the sample of column 1 is split based on saving intentions. LM indicates labor market, Y indicates income; reference category for expectation on Y and LM is increase or stay the same. Reference category for occupational status is pensioner; the unemployed category includes students, housewives, and others. Reference category for making ends meet is very easily.

5. Conclusions

In this paper, we combine a macro and a micro approach to disentangle the different reasons behind the unprecedented drop in private consumption caused by the outbreak of the COVID-19 pandemic and experienced by Italian households. Both aggregate and individual analyses show that spending decisions are only explained in part by the deterioration in the economic fundamentals. The perception of higher income and health risks has a strong impact on spending decisions. The importance of different factors, however, varies by expenditure category and across the household distribution.

This disentangling is also important in order to understand the consequences of the pandemic in the medium run and to interpret actual consumption dynamics in 2021. Italian consumption rebounded with the easing of policy measures and the reduced health risk, consistent with our findings that some transitory factors—the fear of infection above all—played an important role. This is despite the fact that the recovery in household consumption was partial and uneven: goods expenditure has fully recovered pre-pandemic levels, while services still lag behind, as pressures coming from pent-up demand are limited for the latter category (Beraja and Wolf 2021). The saving rate has fallen from the peak levels attained during the recession, but it remains higher than pre-pandemic figures because of a heightened precautionary attitude. Households may prefer to preserve a saving buffer even when the epidemic is fully under control, as long as they remain uncertain about their economic situation and the risk of new pandemic events (Ercolani, Guglielminetti, and Rondinelli 2021).

Looking ahead, our findings suggest that during the ongoing recovery from the COVID-19 pandemic, policymakers could shore up private consumption not only via economic means but also by bringing the epidemic effectively under control and avoiding a widespread surge in the fear of being infected. More generally, precautionary motives and household expectations do play an important role in standard economic recessions as well, possibly slowing down consumption growth if they persist beyond the rebound in income. Moreover, since such factors weigh differently across the economy, this may lead to an uneven recovery, with persistent scarring effects on the hardest-hit sectors and on fragile households.

Appendix A. Macroeconomic Data

Table A.1 summarizes the data used for the analysis in Section 3. In the rest of this section we describe more in detail the main data sources of the macroeconomic time series.

A.1 National Accounts

Consumption data are taken (with the exception of personal care expenditures) from quarterly national accounts—in particular, from the tables reporting the final consumption expenditure of households by expenditure item (COICOP two digit). The two-digit COICOP (Classification of Individual Consumption According to Purpose) categories are 12 (in parentheses we report their share on total consumption in 2019): (i) food and non-alcoholic beverages (14.1 percent); (ii) alcoholic beverages, tobacco, and narcotics (4.1 percent); (iii) clothing and footwear (6.1 percent); (iv) housing, water, electricity, gas, and other fuels (22.6 percent); (v) furnishings, household equipment, and routine household maintenance (6.2 percent); (vi) health (3.5 percent); (vii) transport (12.7 percent); (viii) communications (2.6 percent); (ix) recreation and culture (6.8 percent); (x) education (1.0 percent); (xi) restaurants and hotels (10.2 percent); and (xii) miscellaneous goods and services (10.1 percent). In our analysis by expenditure category, we consider categories (i), (iii), (v), (ix), and (xi), which accounted for more than 43 percent of total household consumption in 2019. Notice that category (iv), which represents a large share of the consumption basket, includes imputed rents of owner-occupied homes.

A.2 Confcommercio

Confcommercio is an Italian confederation of small businesses. Their research group produces a monthly index of real household consumption (ICC), grouped in seven categories: (i) recreation; (ii) hotels and restaurants; (iii) transport; (iv) communication; (v) personal care; (vi) clothing and footwear; and (vii) furnishings, household equipment, and utilities. The total index represents more than 65 percent of total consumption from national accounts excluding imputed rents. In the macro analysis conducted in Section 3 we consider

Table A.1. Description and Sources of Macro Data

Time Series	Description	Source
Total Consumption	Total final consumption expenditure of resident households, chain-linked and seasonally adjusted (s.a.) data	Istat
Consumption: Food and Beverages	Final consumption expenditure of households by expenditure item (COICOP two digit), chain-linked and s.a. data—food and non-alcoholic beverages	Istat
Consumption: Clothing and Footwear	Final consumption expenditure of households by expenditure item (COICOP two digit), chain-linked and s.a. data—clothing and footwear	Istat
Consumption: Furnishings	Final consumption expenditure of households by expenditure item (COICOP two digit), chain-linked and s.a. data—furnishings, household equipment, and routine household maintenance	Istat
Consumption: Recreation and Culture	Final consumption expenditure of households by expenditure item (COICOP two digit), chain-linked and s.a. data—recreation and culture	Istat
Consumption: Restaurants and Hotels	Final consumption expenditure of households by expenditure item (COICOP two digit), chain-linked and s.a. data—restaurants and hotels	Istat
Consumption: Personal Care	Index of real household expenditure in personal care items, s.a. data	Concommercio
Permanent Income	Four-quarters moving average of real gross disposable income, s.a. data	Our elaborations on Istat data
Income Volatility	Rolling standard deviation of the last nine terms of the q-o-q growth rates of the s.a. real gross disposable income	Our elaborations on Istat data
Hours Worked	Total hours worked from quarterly national accounts, s.a. data	Istat
Real Wealth	Housing wealth	Bank of Italy estimates
Financial Wealth	Financial activities held by households from financial accounts	Bank of Italy
Interest Rate	Long-term interest rate, 10 years maturity, new business coverage	ECB
Unemployment	Unemployment expectations from monthly consumer surveys, balances, s.a. data, quarterly averages	Istat
Expectations	Principal component of the dispersion of the answers to three questions of consumer surveys, regarding expectations on the personal and general economic situation and on unemployment, quarterly averages	Istat
Disagreement		

the indicator of personal care expenditures, as this category is not available in the COICOP 2 classification of the national accounts.

A.3 Consumer Surveys

The European Commission conducts harmonized business and consumer in the European Union (EU) and in the applicant countries at monthly frequency. Surveys are carried out at the national level by the national statistical institutes. For the purpose of this paper, we take the data from Istat (whose seasonally adjusted balances of responses can slightly differ from those taken from the European Commission due to a different treatment of the raw data) and exploit three questions included in the monthly consumer surveys. The questions are (in Italian in the original questionnaire):

1. (*Expectations on general economic situation*) In your opinion, in the next 12 months, the general economic situation of Italy:
 - (a) Will markedly improve
 - (b) Will slightly improve
 - (c) Will remain unchanged
 - (d) Will slightly worsen
 - (e) Will markedly worsen
 - (f) I don't know
2. (*Expectations on personal economic situation*) In your opinion, over the next 12 months, the economic situation of your family:
 - (a) Will markedly improve
 - (b) Will slightly improve
 - (c) Will remain unchanged
 - (d) Will slightly worsen


- (e) Will markedly worsen
 - (f) I don't know
3. (*Unemployment expectations*) In your opinion, over the next 12 months, the number of unemployed in Italy:
- (a) Will strongly increase
 - (b) Will slightly increase
 - (c) Will remain unchanged
 - (d) Will slightly decrease
 - (e) Will strongly decrease
 - (f) I don't know

From the three questions reported above, we exploit two measures that we include in the macro regression (1). First, we use the seasonally adjusted balances of responses regarding unemployment expectations $((a+b)-(d+e))$, as provided directly by Istat. Second, for each of the three above-mentioned questions we build a measure of dispersion of responses (disagreement). Assume that F_j is the cumulated frequency of the answers up to j , starting from the most negative ones (hence from (e) to (a)). Our measure of dispersion in the answers to question x is

$$d_x = \sum F_e(1 - F_e) + F_d(1 - F_d) + F_c(1 - F_c) + F_b(1 - F_b) + F_a(1 - F_a).$$

This indicator takes values between 0 (no dispersion) and 1 (maximum dispersion). We then standardize the disagreement indicators d_1 , d_2 , and d_3 (we subtract the mean and divide by the standard deviation) and take the first principal component. This is the measure of disagreement we include in the macro regression (1).

Figure B.1. Unexplained Consumption Growth in Macro Regressions (realized – predicted)



Source: Our computations on data from Istat, Bank of Italy, ECB, and Confindustria.

Note: For each consumption item we estimate the regression (1) and compute the difference between realized and predicted y-o-y consumption growth (residual).

Appendix C. The Bank of Italy’s Special Survey of Italian Households (SSIH) Questionnaire

In 2020, because of the spread of the COVID-19 epidemic, the Bank of Italy suspended the fieldwork activities of the surveys on households carried out through personal interviews, in order to ensure the health of both households and interviewers. To fill the information gap and gather timely information on the economic situation of households during the pandemic crisis, the Bank of Italy launched a Special Survey of Italian Households, administered using remote devices.

In the first edition of the survey, carried out between the end of April and the start of May 2020, two other survey techniques (namely interviews over the phone and the web) were used in addition to the touchscreen device, which was the only method used for the second, third, and fourth editions. The main findings and the methodology of SSIH are described in Neri and Zanichelli (2020)

and in the box “Italian Households’ Assessments and Expectations during the Current Public Health Emergency” (Bank of Italy 2020b).

The Survey is representative of the Italian population. The system of sample weights was calibrated to the information provided by the National Institute of Statistics (Istat) on gender, age group, geographical area (North, Center, South, and Islands), and degree of education and employment status (employee, self-employed, retired, unemployed, other) of the Italian population aged 18 and over.

C.1 The COVID-19 Module in Wave 3 of SSIH

The third wave of the survey was conducted at the end of November 2020, and the total sample amounted to more than 2,000 households, which had also been interviewed in the second edition of the survey. The main findings and the methodology used are described in Rondinelli and Zanichelli (2021a) and in the box “Italian Households during the Epidemic: The Bank of Italy’s Survey” (Bank of Italy 2021a).

1. Compared to before the COVID-19 pandemic, in the last month how frequently you did these activities?

1= never done or much less often; 2= less often; 3= with the same frequency, 4= often, 5= more often

	1	2	3	4	5
1. shop for food and other essential goods in stores					
2. make purchases in stores of clothing, footwear, etc.					
3. go to the hairdresser, beautician and other personal services					
4. go out to shop for furniture, appliances, etc.					
5. go to hotels/bars/restaurants					

2. (If you answered 1 or 2 to at least one of 1.2, 1.3, and 1.5) Why did you make some expenses less often?
- Lower financial resources
 - Other reasons

3. (If “Other reasons” at previous question) Distribute 100 points between these three alternatives, based on what you think are most likely: give a high score to those you think are most likely, a low one to those you deem least likely.
- Given the containment measures, some purchases were hindered/impossible
 - Fears of contagion
 - Put money aside for unexpected events
4. Compared to normal times, how frequently would you carry out the aforementioned activities if the daily number of infections in your region were (a) less than 10, (b) between 10 and 100, (c) between 100 and 1000, (d) greater than 1000? For each of the cases (a), (b), (c), (d) choose between 1= never done or much less often; 2= less often; 3= with the same frequency, 4= often, 5= more often

	1	2	3	4	5
1. shop for food and other essential goods in stores					
2. make purchases in stores of clothing, footwear, etc.					
3. go to the hairdresser, beautician and other personal services					
4. go out to shop for furniture, appliances, etc.					
5. go to hotels/bars/restaurants					

C.2 The COVID-19 Module in Wave 2 of SSIH

The second wave of the survey was conducted from late August to early September 2020; the total sample amounted to more than 2,300 households. The main findings and the methodology are described in Rondinelli and Zanichelli (2020) and in the box “Italian Households during the Epidemic: The Bank of Italy’s Survey” (Bank of Italy 2020a).

1 As in the COVID-19 module of the wave 3 of SSIH.

2 Not included

3 Not included

4 As in the COVID-19 module of the wave 3 of SSIH.

C.3 The COVID-19 Module in Wave 4 of SSIH

The fourth wave of the survey was conducted between February and March 2021, and the total sample amounted to more than 2,000 households; almost 1,800 had also been interviewed in the third edition of the survey. The main findings and the methodology used are described in Rondinelli and Zanichelli (2021b) and in the box “Italian Households during the Epidemic: The Bank of Italy’s Survey” (Bank of Italy 2021b).

1 As in the COVID-19 module of the wave 3 of SSIH.

2+3 (If you answered “1” or “2” to at least one of 1.2, 1.3, and 1.5 of the COVID-19 module) Why did you spend some expenses less often? Distribute 100 points between these three alternatives, based on what you think are most likely: give a high score to those you think are most likely, a low one to those it deems least likely.

- Lower financial resources
- Given the containment measures, some purchases were hindered/impossible
- Fears of contagion
- Put money aside for unexpected events

4 Not included

Accumulated Saving: Did your household spend less than your annual income in 2020, putting aside some savings? Yes/No

C.4 Construction of the Relevant Variables

Expected Spending: Consider the expenditure for food, clothing and footwear, home goods and services. How will your household change total expenditure for these items in the next 3 months? It will increase; stay the same; decrease

Expected Drop in Spending: (If expected spending “will decrease”) By how much? Less than 10%; Between 10% and 20%; Between 20% and 30%; More than 30%

Saving Intentions: In the next 12 months, do you expect to: spend less than the entire yearly income and succeed in saving; spend the entire yearly income and not to manage to save anything; spend more than the entire yearly income, drawing on savings or borrowing

Making Ends Meet: Before the start of the pandemic, your household income allowed you to reach the end of the month with: great difficulty, difficulty, some difficulty, easily enough, easily, very easily

Expectation on Y: How do you expect your household income to change in 2021 compared to 2020? It will increase, stay the same, decrease

Expectation on LM: In your opinion, how will the labor market situation in Italy develop over the next 12 months? It will: considerably improve, slightly improve, stay the same, slightly deteriorate, considerably deteriorate, don’t know

FEAR: Dummy variable equal to one whether the households answered “1” to questions 1.2, 1.3, and 1.5 of the COVID-19 module

FEAR CASES: Dummy variable equal to one whether the households answered “1” or “2” to question 4.5(a) of the COVID-19 module, i.e., if households decided to eliminate or strongly reduce their spending in hotels, bars, and restaurants even when in the region of residence there were fewer than 10 new cases of COVID-19 per day.³⁸

Uncertainty: Dummy variable equal to one if the self-reported probability of (i) losing the job is higher than 25% for employed household members and (ii) finding a job is lower than 25% for unemployed individuals.

(i) (If “Employee” or “Self-employed”) What is the probability that you will lose your job over the next 12 months? (Answer between a minimum of “0” if you are sure you are working and a maximum of “100” if you are sure you are not working.)

(ii) (if “Unemployed”) What is the probability that you will find a (new) job over the next 12 months? (Answer between a minimum

³⁸Despite the fact that the reference to the number of cases is in absolute terms, 10 can be considered a sufficiently small number, even in smaller regions.

of “0” if you are sure you are not working and a maximum of “100” if you are sure you are working.)

Red/Orange/Yellow Zone: Dummy variables equal to one for regions in red/orange/yellow areas at the moment of the interview. Starting from the beginning of November 2020, Italian regions were differentiated on a weekly basis according to the degree of restrictions in place, identified either by a yellow, orange, or red color (with increasing severity). When wave 3 of SSIH began, the red zones were the Autonomous Province of Bolzano and the regions of Calabria, Campania, Lombardy, Piedmont, Tuscany, and Valle d’Aosta; the orange zones were Abruzzo, Basilicata, Emilia-Romagna, Friuli Venezia Giulia, Liguria, Marche, Puglia, Sicily, and Umbria. All the other regions and the Autonomous Province of Trento were designated yellow zones (Prime Minister’s Decree of 3 November 2020 and subsequent ordinances).

Government Measures: A dummy variable equal to one for households having received fiscal support to income by the Italian government.

Table C.3. SSIH Data: Summary Statistics

Variable	Mean	Std. Dev.	Obs.
20–34	0.07	0.255	2,077
35–54	0.368	0.482	2,077
55–64	0.208	0.406	2,077
65+	0.355	0.478	2,077
None or Primary	0.191	0.393	2,077
Middle School	0.373	0.484	2,077
High School	0.295	0.456	2,077
College and Beyond	0.141	0.348	2,077
Male	0.684	0.465	2,077
Female	0.316	0.465	2,077
North	0.489	0.5	2,077
Center	0.203	0.402	2,077
South	0.308	0.462	2,077
Employee	0.403	0.491	2,077
Self-Employed	0.092	0.288	2,077
Unemployed	0.14	0.347	2,077
Pensioner	0.365	0.482	2,077
Making Ends Meet (Great Difficulty)	0.188	0.391	2,077
Making Ends Meet (Difficulty)	0.351	0.477	2,077
Making Ends Meet (Easily)	0.305	0.46	2,077
Making Ends Meet (Very Easily)	0.156	0.363	2,077
Expectation on Y (Increase)	0.146	0.353	2,077
Expectation on Y (Stable)	0.658	0.475	2,077
Expectation on Y (Decrease)	0.197	0.398	2,077
Expectation on LM (Increase)	0.207	0.405	2,077
Expectation on LM (Stable)	0.147	0.354	2,077
Expectation on LM (Decrease)	0.583	0.493	2,077
Expectation on C Less Y in 12 Months ($S = 1$)	0.419	0.493	2,077
Expectation on C in 3 Months (Decrease, $ncexp = 1$)	0.328	0.47	2,077
Expected C Will Decrease < 10%	0.146	0.353	675
Expected C Will Decrease 10–20%	0.348	0.477	675
Expected C Will Decrease 20–30%	0.208	0.406	675
Expected C Will Decrease > 30%	0.297	0.457	675
<i>Survey Questions on the Reasons for Cutting Expenditure</i>			
Lower Financial Resources; % of Households	0.464	0.499	1,838
Restrictive Measures; Points Assigned Mean (Median)	31.64 (30)	22.537	1,032
Fears of Contagion; Points Assigned Mean (Median)	39.016 (40)	23.292	1,032
Unexpected Events; Points Assigned Mean (Median)	29.344 (25)	21.007	1,032

(continued)

Table C.3. (Continued)

Variable	Mean	Std. Dev.	Obs.
<i>Built-in Questions on the Reasons for Cutting Expenditure</i>			
FEAR	0.277	0.448	2,077
FEARS CASE	0.58	0.494	2,077
Uncertainty	0.316	0.465	1,408
Red Zone	0.405	0.491	2,077
Orange Zone	0.37	0.483	2,077
Yellow Zone	0.225	0.417	2,077
Note: Our calculations from wave 3 of SSIH, using sample weights. See Appendix C for the variable construction. The unemployed category includes students, housewives, and others. LM indicates labor market, Y indicates income, C indicates expenditure, S indicates saving.			

Appendix D. Microeconomic Evidence: Additional Results

Table D.1. Determinants of the Reasons for Cutting Expenditure (multinomial logit)

	Economic	Fears	Deal with Unexpected Events
Age	−0.029 (0.049)	−0.047 (0.051)	−0.021 (0.049)
Age Sq.	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Employee	−0.107 (0.349)	−0.702** (0.356)	−0.394 (0.342)
Self-Employed	0.268 (0.364)	−0.729* (0.387)	−0.398 (0.359)
Unemployed	−0.023 (0.422)	−0.550 (0.453)	−0.484 (0.430)
Making Ends Meet (Great Difficulty)	1.814*** (0.292)	−0.010 (0.319)	0.035 (0.300)
Making Ends Meet (Difficulty)	1.587*** (0.244)	0.092 (0.243)	0.503** (0.235)
Making Ends Meet (Easily)	0.309 (0.223)	−0.239 (0.214)	0.166 (0.208)
Expectation on Y (Decrease)	0.354* (0.211)	−0.016 (0.244)	−0.052 (0.216)
Expectation on LM (Decrease)	0.259 (0.167)	0.053 (0.180)	−0.088 (0.163)
Red Zone	0.011 (0.194)	−0.126 (0.208)	−0.089 (0.191)
Orange Zone	0.361* (0.207)	0.053 (0.224)	0.226 (0.207)
FEAR	0.765*** (0.199)	0.566*** (0.210)	0.502** (0.200)
Uncertainty	0.841*** (0.212)	0.371* (0.223)	0.499** (0.212)
Constant	0.128 (1.265)	2.017 (1.353)	1.239 (1.252)

Note: ** $p < 0.05$, *** $p < 0.01$. Number of observations: 1,238. Reference for the multinomial model is “restrictive measures.” Our calculations from wave 3 of SSIH. LM indicates labor market, Y indicates income; reference category for expectation on Y and LM is increase or stay the same. Reference category for occupational status is pensioner; the unemployed category includes students, housewives, and others. Reference category for making ends meet is very easily.

Table D.2. Channels for Cutting Expenditure
(different waves, linear models)

	Wave 3	Wave 2	Waves 2 and 3	Wave 4
	(1)	(2)	(3)	(4)
Age	0.010 (0.006)	0.014** (0.006)	0.011** (0.004)	0.013*** (0.004)
Age Sq.	-0.000 (0.000)	-0.000* (0.000)	-0.000** (0.000)	-0.000*** (0.000)
Employee	0.082 (0.055)	0.001 (0.048)	0.042 (0.036)	-0.016 (0.031)
Self-Employed	0.091 (0.058)	0.074 (0.051)	0.081** (0.038)	-0.004 (0.036)
Unemployed	0.076 (0.061)	0.051 (0.051)	0.067* (0.040)	-0.009 (0.036)
Making Ends Meet (Great Difficulty)	0.073* (0.040)	0.179*** (0.038)	0.126*** (0.027)	0.025 (0.030)
Making Ends Meet (Difficulty)	0.101*** (0.035)	0.067* (0.035)	0.083*** (0.025)	0.060** (0.026)
Making Ends Meet (Easily)	0.010 (0.035)	0.033 (0.035)	0.019 (0.025)	0.008 (0.026)
Expectation on Y (Decrease)	0.286*** (0.029)	0.263*** (0.026)	0.274*** (0.019)	0.176*** (0.023)
Expectation on LM (Decrease)	0.064*** (0.024)	0.097*** (0.023)	0.082*** (0.016)	0.094*** (0.017)
Red Zone	0.044 (0.030)	-0.016 (0.028)	0.014 (0.020)	0.025 (0.018)
Orange Zone	0.044 (0.030)	-0.070** (0.029)	-0.011 (0.021)	-0.102 (0.071)
FEAR	0.228*** (0.027)	0.145*** (0.026)	0.186*** (0.019)	0.213*** (0.021)
Uncertainty	0.115*** (0.028)	0.082*** (0.026)	0.098*** (0.019)	0.111*** (0.020)
Constant	-0.312* (0.163)	-0.334** (0.157)	-0.309*** (0.113)	-0.275*** (0.106)
Obs.	1,408	1,512	2,920	2,400
R ²	0.205	0.193	0.191	0.140

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Our calculations from wave 2 (April–May 2020), wave 3 (November 2020), and wave 4 (February–March 2021) of SSIH (see Appendix C). Dependent variable equal to one for households expecting a drop in expenditure in the next three months. LM indicates labor market, Y indicates income; reference category for expectation on Y and LM is increase or stay the same. Reference category for occupational status is pensioner; the unemployed category includes students, housewives, and others. Reference category for making ends meet is very easily.

Table D.3. Channels for Cutting Expenditure by Job Status (linear models)

	Baseline	Employed	Self-Employed	Unemployed	Pensioners
	(1)	(2)	(3)	(4)	(5)
Age	0.013** (0.006)	0.001 (0.009)	-0.011 (0.020)	0.044*** (0.016)	-0.029 (0.072)
Age Sq.	-0.000* (0.000)	0.000 (0.000)	0.000 (0.000)	-0.000*** (0.000)	0.000 (0.001)
Making Ends Meet (Great Difficulty)	0.074* (0.040)	0.104** (0.049)	-0.130 (0.116)	0.193* (0.113)	-0.143 (0.124)
Making Ends Meet (Difficulty)	0.104*** (0.035)	0.067 (0.043)	-0.067 (0.100)	0.242** (0.113)	0.453*** (0.108)
Making Ends Meet (Easily)	0.010 (0.035)	0.003 (0.042)	-0.068 (0.101)	0.065 (0.149)	0.073 (0.096)
Expectation on Y (Decrease)	0.284*** (0.029)	0.264*** (0.038)	0.268*** (0.070)	0.317*** (0.079)	0.307*** (0.080)
Expectation on LM (Decrease)	0.064*** (0.024)	0.067** (0.029)	0.080 (0.066)	0.038 (0.082)	-0.019 (0.073)
Red Zone	0.044 (0.030)	0.037 (0.036)	0.015 (0.082)	0.165* (0.097)	0.107 (0.090)
Orange Zone	0.043 (0.030)	0.074* (0.038)	-0.083 (0.079)	-0.051 (0.096)	0.212** (0.091)
FEAR	0.228*** (0.027)	0.186*** (0.035)	0.339*** (0.071)	0.265*** (0.075)	-0.002 (0.087)
Uncertainty	0.117*** (0.026)	0.157*** (0.035)	0.195*** (0.069)	0.077 (0.094)	-0.297*** (0.092)
Constant	-0.289* (0.162)	-0.038 (0.232)	0.579 (0.518)	-1.197*** (0.416)	1.013 (2.528)
Obs.	1,408	904	212	143	131
R ²	0.204	0.181	0.287	0.394	0.344
Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Our calculations from wave 3 of SSIH. Dependent variable equal to one for households expecting a drop in expenditure in the next three months. LM indicates labor market, Y indicates income; reference category for expectation on Y and LM is increase or stay the same. Reference category for making ends meet is very easily. The unemployed category does not include students and housewives.					

Table D.4. Channels for Cutting Expenditure by Age (linear models)

	Baseline	<35	35–54	55–65	>65
	(1)	(2)	(3)	(4)	(5)
Age	0.010 (0.006)	1.092*** (0.375)	–0.036 (0.049)	0.483* (0.291)	–0.397*** (0.164)
Age Sq.	–0.000 (0.000)	–0.018*** (0.006)	0.000 (0.001)	–0.004* (0.002)	0.003*** (0.001)
Employee	0.082 (0.055)	0.359* (0.188)	0.160 (0.339)	–0.036 (0.090)	0.055 (0.100)
Self-Employed	0.091 (0.058)	0.493*** (0.215)	0.145 (0.341)	0.013 (0.099)	–0.020 (0.096)
Unemployed	0.076 (0.061)	0.000 (.)	0.131 (0.342)	0.066 (0.105)	–0.015 (0.123)
Making Ends Meet (Great Difficulty)	0.073* (0.040)	0.214 (0.160)	0.058 (0.052)	0.027 (0.080)	0.069 (0.126)
Making Ends Meet (Difficulty)	0.101*** (0.035)	0.080 (0.136)	0.063 (0.046)	0.079 (0.074)	0.178* (0.106)
Making Ends Meet (Easily)	0.010 (0.035)	–0.063 (0.135)	0.013 (0.046)	0.061 (0.076)	–0.100 (0.102)

(continued)

Table D.4. (Continued)


	Baseline	<35	35–54	55–65	>65
	(1)	(2)	(3)	(4)	(5)
Expectation on Y (Decrease)	0.286*** (0.029)	0.042 (0.121)	0.330*** (0.038)	0.282*** (0.059)	0.319*** (0.086)
Expectation on LM (Decrease)	0.064*** (0.024)	0.093 (0.092)	0.078** (0.031)	0.085* (0.048)	–0.055 (0.073)
Red Zone	0.044 (0.030)	–0.206* (0.114)	0.059 (0.039)	0.117** (0.058)	0.117 (0.093)
Orange Zone	0.044 (0.030)	–0.078 (0.117)	0.069* (0.040)	0.082 (0.059)	0.094 (0.097)
FEAR	0.228*** (0.027)	0.312** (0.121)	0.260*** (0.035)	0.185*** (0.054)	0.115 (0.084)
Uncertainty	0.115*** (0.028)	0.348*** (0.125)	0.099*** (0.036)	0.102* (0.054)	0.016 (0.091)
Constant	–0.312* (0.163)	–16.385*** (5.616)	0.461 (1.129)	–13.700 (8.682)	14.959** (6.174)
Obs.	1,408	88	775	386	159
R ²	0.205	0.370	0.252	0.230	0.246
Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Our calculations from wave 3 of SSIH. Dependent variable equal to one for households expecting a drop in expenditure in the next three months. LM indicates labor market, Y indicates income; reference category for expectation on Y and LM is increase or stay the same. Reference category for making ends meet is very easily. The unemployed category includes students, housewives, and others.					

**Table D.5. Connecting the Drop in Consumption
with Saving in 2020 (linear models)**

	Baseline	Accumulate Savings in 2020	Did Not Accumulate Savings in 2020
	(1)	(2)	(3)
Age	0.010 (0.006)	−0.003 (0.010)	0.017* (0.009)
Age Sq.	−0.000 (0.000)	0.000 (0.000)	−0.000 (0.000)
Employee	0.082 (0.055)	0.047 (0.082)	0.083 (0.081)
Self-Employed	0.091 (0.058)	0.067 (0.087)	0.074 (0.083)
Unemployed	0.076 (0.061)	0.054 (0.103)	0.069 (0.085)
Making Ends Meet (Great Difficulty)	0.073* (0.040)	0.041 (0.073)	−0.022 (0.070)
Making Ends Meet (Difficulty)	0.101*** (0.035)	0.041 (0.050)	0.052 (0.067)
Making Ends Meet (Easily)	0.010 (0.035)	−0.036 (0.043)	0.024 (0.070)
Expectation on Y (Decrease)	0.286*** (0.029)	0.276*** (0.053)	0.281*** (0.037)
Expectation on LM (Decrease)	0.064*** (0.024)	0.067* (0.036)	0.052 (0.035)
Red Zone	0.044 (0.030)	0.023 (0.045)	0.030 (0.043)
Orange Zone	0.044 (0.030)	0.007 (0.048)	0.066 (0.043)
FEAR	0.228*** (0.027)	0.156*** (0.043)	0.274*** (0.039)
Uncertainty	0.115*** (0.028)	0.072 (0.047)	0.132*** (0.037)
Constant	−0.312* (0.163)	0.111 (0.252)	−0.442* (0.233)
Obs.	1,408	514	751
R ²	0.205	0.128	0.217

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Our calculations from wave 3 of SSIH considering households interviewed also in wave 4 (see Appendix C) and replying to the question on the savings accumulated in 2020. Dependent variable equal to one for households expecting a drop in expenditure in the next three months for column 1. Dependent variable equal to one for households (i) accumulating savings in 2020 for column 2 and (ii) non-accumulating savings in 2020 for column 3. LM indicates labor market, Y indicates income; reference category for expectation on Y and LM is increase or stay the same. Reference category for making ends meet is very easily. The unemployed category includes students, housewives, and others.

Figure D.1. Contribution of the Reasons for Cutting Expenditure by Age



Note: Our calculations from wave 3 of SSIH. Contributions are computed at the sample means of each age category.

References

- Alexander, D., and E. Karger. 2023. “Do Stay-at-Home Orders Cause People to Stay at Home? Effects of Stay-at-Home Orders on Consumer Behavior.” *Review of Economics and Statistics* 105 (4): 1017–27. ISSN 0034-6535. doi: 10.1162/rest_a_01108. https://doi.org/10.1162/rest_a_01108.
- Andersen, A. L., E. T. Hansen, N. Johannesen, and A. Sheridan. 2020. “Pandemic, Shutdown and Consumer Spending: Lessons from Scandinavian Policy Responses to Covid-19.” *Papers* 2005.04630, arXiv.org, May. <https://ideas.repec.org/p/arx/papers/2005.04630.html>.
- Ando, A., and F. Modigliani. 1963. “The ‘Life Cycle’ Hypothesis of Saving: Aggregate Implications and Tests.” *American Economic Review* 53 (5): 55–84.
- Ardizzi, G., A. Nobili, and G. Rocco. 2020. “A Game Changer in Payment Habits: Evidence from Daily Data during a Pandemic.” *Questioni di Economia e Finanza (Occasional Papers)*

- No. 591, Bank of Italy. https://ideas.repec.org/p/bdi/opques/gef_591.20.html.
- Bachmann, R., B. Born, O. Goldfayn-Frank, G. Kocharkov, R. Luetticke, and M. Weber. 2021. "A Temporary VAT Cut as Unconventional Fiscal Policy." NBER Working Paper No. 29442. <http://www.nber.org/papers/w29442>.
- Baker, S. R., R. A. Farrokhnia, S. Meyer, M. Pagel, and C. Yannelis. 2020. "How Does Household Spending Respond to an Epidemic? Consumption during the 2020 COVID-19 Pandemic." *Review of Asset Pricing Studies* 10 (4): 834–62. ISSN 2045-9920. doi: 10.1093/rapstu/raaa009. <https://doi.org/10.1093/rapstu/raaa009>.
- Bank of Italy. 2020a. "Economic Bulletin." Technical Report No. 4.
- . 2020b. "Annual Report for 2019." Technical Report, May.
- . 2021a. "Economic Bulletin." Technical Report No. 1.
- . 2021b. "Economic Bulletin." Technical Report No. 2.
- . 2021c. "Annual Report for 2020." Technical Report, May.
- Beraja, M., and C. Wolf. 2021. "Demand Composition and the Strength of Recoveries." NBER Working Paper No. 29304.
- Chenarides, L., C. Grebitus, J. L. Lusk, and I. Printezis. 2021. "Food Consumption Behavior during the Covid-19 Pandemic." *Agribusiness* 37 (1): 44–81. doi: <https://doi.org/10.1002/agr.21679>. <https://onlinelibrary.wiley.com/doi/abs/10.1002/agr.21679>.
- Chetty, R., J. N. Friedman, N. Hendren, M. Stepner, and the Opportunity Insights Team. 2020. "The Economic Impacts of Covid-19: Evidence from a New Public Database Built Using Private Sector Data." NBER Working Paper No. 27431. <http://www.nber.org/papers/w27431>.
- Christelis, D., D. Georgarakos, T. Jappelli, and G. Kenny. 2020a. "The Covid-19 Crisis and Consumption: Survey Evidence from Six EU Countries." ECB Working Paper No. 2507. <https://ideas.repec.org/p/ecb/ecbwps/20202507.html>.
- Christelis, D., D. Georgarakos, T. Jappelli, and M. van Rooij. 2020b. "Consumption Uncertainty and Precautionary Saving." *Review of Economics and Statistics* 102 (1): 148–61, March. <https://ideas.repec.org/a/tpr/restat/v102y2020i1p148-161.html>.

- Coibion, O., D. Georgarakos, Y. Gorodnichenko, G. Kenny, and M. Weber. 2021. "The Effect of Macroeconomic Uncertainty on Household Spending." NBER Working Paper No. 28625. <http://www.nber.org/papers/w28625>.
- Coibion, O., Y. Gorodnichenko, and M. Weber. 2020a. "How Did U.S. Consumers Use Their Stimulus Payments?" NBER Working Paper No. 27693. <http://www.nber.org/papers/w27693>.
- . 2020b. "The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending." *Covid Economics: Vetted and Real-Time Papers* (CEPR) 20 (May).
- Cox, N., P. Ganong, P. J. Noel, J. S. Vavra, A. Wong, D. Farrell, and F. E. Greig. 2020. "Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Savings Data." NBER Working Paper No. 27617. <https://ideas.repec.org/p/nbr/nberwo/27617.html>.
- D'Acunto, F., D. Hoang, M. Paloviita, and M. Weber. 2021. "Effective Policy Communication: Targets versus Instruments." Technical Report No. 147, Karlsruhe Institute of Technology (KIT). <https://ideas.repec.org/p/zbw/kitwps/147.html>.
- De Bonis, R., D. Liberati, J. Muellbauer, and C. Rondinelli. 2020. "Consumption and Wealth: New Evidence from Italy." Temi di discussione (Economic Working Paper) No. 1304, Bank of Italy. <https://ideas.repec.org/p/bdi/wptemi/td.1304.20.html>.
- Dossche, M., and S. Zlatanos. 2020. "COVID-19 and the Increase in Household Savings: Precautionary or Forced?" *Economic Bulletin Boxes* (European Central Bank), Vol. 6. <https://ideas.repec.org/a/ecb/ecbbox/202000065.html>.
- Emiliozzi, S., C. Rondinelli, and S. Villa. 2023. "Consumption during the Covid-19 Pandemic: Evidence from Italian Credit Cards." Occasional Paper No. 769, Bank of Italy. <https://ideas.repec.org/p/bdi/opques/qef.769.23.html>.
- Ercolani, V., E. Guglielminetti, and C. Rondinelli. 2021. "Fears of Future: The Saving Dynamics after the Covid-19 Outbreak." Covid-19 Note, Bank of Italy. https://www.bancaditalia.it/pubblicazioni/note-covid-19/2021/Note_saving_EGR_14giugno2021.pdf.
- Ganong, P., and P. Noel. 2019. "Consumer Spending during Unemployment: Positive and Normative Implications." *American Economic Review* 109 (7): 2383–2424. doi: 10.1257/

- aer.20170537. <https://www.aeaweb.org/articles?id=10.1257/aer.20170537>.
- Guiso, L., M. Paiella, and I. Visco. 2006. "Do Capital Gains Affect Consumption? Estimates of Wealth Effects from Italian Households' Behavior." In *Long-run Growth and Short-run Stabilization: Essays in Memory of Albert Ando*, ed. L. R. Klein. Edward Elgar Publishing.
- Hodbod, A., C. Hommes, S. J. Huber, and I. Salle. 2021. "The COVID-19 Consumption Game-changer: Evidence from a Large-scale Multi-country Survey." *European Economic Review* 140 (November): Article 103953. ISSN 0014-2921. doi: <https://doi.org/10.1016/j.euroecorev.2021.103953>. <https://www.sciencedirect.com/science/article/pii/S0014292121002440>.
- Immordino, G., T. Jappelli, T. Oliviero, and A. Zazzaro. 2022. "Fear of Covid-19 Contagion and Consumption: Evidence from a Survey of Italian Households." *Health Economics* 31 (3): 496–507. doi: <https://doi.org/10.1002/hec.4464>. <https://onlinelibrary.wiley.com/doi/abs/10.1002/hec.4464>.
- Istat. 2021. "Le spese per i consumi delle famiglie: anno 2020." Technical Report, Istituto Nazionale di Statistica. https://www.istat.it/it/files/2021/06/REPORT_CONSUMI_FAMIGLIE_2020.pdf.
- Jordà, Ò., S. R. Singh, and A. M. Taylor. 2022. "Longer-Run Economic Consequences of Pandemics." *Review of Economics and Statistics* 104 (1): 166–75. ISSN 0034-6535. doi: 10.1162/rest_a_01042. https://doi.org/10.1162/rest_a_01042.
- Neri, A., and F. Zanichelli. 2020. "The Main Results of the Special Survey of Italian Households in 2020." Covid-19 Note, Bank of Italy, June. <https://www.bancaditalia.it/pubblicazioni/note-Covid-19/2020/Evi-preliminari-ind-straord-famiglie.pdf>.
- Paiella, M. 2007. "Does Wealth Affect Consumption? Evidence for Italy." *Journal of Macroeconomics* 29 (1, March): 189–205. <https://ideas.repec.org/a/eee/jmacro/v29y2007i1p189-205.html>.
- Paiella, M., and L. Pistaferri. 2017. "Decomposing the Wealth Effect on Consumption." *Review of Economics and Statistics* 99 (4): 710–21. <https://EconPapers.repec.org/RePEc:trp:restat:v:99:y:2017:i:4:p:710-721>.

- Perez-Quiros, G., and A. Timmermann. 2000. "Firm Size and Cyclical Variations in Stock Returns." *Journal of Finance* 55 (3, June): 1229–62.
- Rodano, L., and C. Rondinelli. 2014. "The Italian Household Consumption: A Comparison Among Recessions." *Politica Economica* (2–3): 203–34. <https://ideas.repec.org/a/mul/je8794/doi10.1429-80193y2014i2-3p203-234.html>.
- Rondinelli, C., and F. Zanichelli. 2020. "The Main Results of the Second Wave of the Special Survey of Italian Households in 2020." Covid-19 Note, Bank of Italy, November. <https://www.bancaditalia.it/pubblicazioni/note-Covid-19/2020/Nota-covid-2020.11.19.pdf>.
- . 2021a. "The Main Results of the Third Wave of the Special Survey of Italian Households in 2020." Covid-19 Note, Bank of Italy, March. https://www.bancaditalia.it/pubblicazioni/note-covid-19/2021/2021.03.30_Nota_Covid_principali_risultati_3a_ed_ind_straord_famiglie_italiane_2020.pdf.
- . 2021b. "The Main Results of the Fourth Wave of the Special Survey of Italian Households." Covid-19 Note, Bank of Italy, May. https://www.bancaditalia.it/pubblicazioni/note-covid-19/2021/en_Nota_Covid_ISF4_210521.pdf?language_id=1.

Anchored or Not: How Much Information Does 21st Century Data Contain on Inflation Dynamics?*

Michael T. Kiley
Federal Reserve Board

Inflation was low and stable in the United States during the first two decades of the 21st century and broke out of its stable range in 2021. Experience in the early 21st century differed from that of the second half of the 20th century, when inflation showed persistent movements including the “Great Inflation” of the 1970s. This analysis examines the extent to which the experience from 2000–19 should lead a Bayesian decisionmaker to update their assessment of inflation dynamics. Given a prior for inflation dynamics consistent with 1960–99 data, a Bayesian decisionmaker would not update their view of inflation persistence in light of 2000–19 data unless they placed very low weight on their prior information. In other words, 21st century data contain very little information to dissuade a Bayesian decisionmaker of the view that inflation fluctuations are persistent, or “unanchored.” The intuition for, and implications of, this finding are discussed.

JEL Codes: E31, C11, E50.

1. Introduction

Consumer price inflation in the United States, as measured by the consumer price index, jumped to just above 7 percent in the 12 months ending in December 2021. Inflation in 2021 reached the highest level seen since the early 1980s. The jump in inflation

*Michael Kiley is Deputy Director of the Division of Financial Stability of the Federal Reserve Board. The views expressed herein are those of the author, and do not reflect those of the Federal Reserve Board or its staff. This research has benefited from comments at the CEF 2022 conference in Dallas and the CEBRA 2022 conference in Barcelona. E-mail: mkiley@frb.gov.

outside of the range experienced over several decades has raised questions regarding the speed with which, or the degree to which, inflation may return to the 2 percent range consistent with the Federal Reserve's inflation objective. The answers to these questions hinge on the nature of the inflation process, including its persistence and the impact of economic slack on the inflation outlook.

A critical consideration in any work pursuing these questions is how to weigh data from different time periods. For example, one approach would involve specifying a (set of) model(s) and examining the stability of the inflation process across time periods as indicated by statistical tests. An alternative approach would involve a (set of) model(s) in which the parameters of the inflation process evolve over time (a time-varying parameter model) and would combine the model and data to estimate the evolution of parameters and resulting implications for the inflation outlook. Substantial bodies of research have considered both approaches.

This research pursues a different tack. The approach herein is the one a Bayesian decisionmaker would follow. The decisionmaker is endowed with a prior regarding the inflation process consistent with observed U.S. inflation data over the second half of the 20th century. The analysis examines how the Bayesian decisionmaker would use data since 2000 to update his prior view. The approach of the Bayesian decisionmaker has several advantages relative to other approaches. The first advantage is that the approach has not been used to inform an assessment of inflation dynamics and hence provides a new (and different) perspective. The second advantage is that the approach is a natural way to combine prior information/experience with recent data, in two senses: the approach flows directly from the standard approach to combining prior information with subsequent data laid out by Thomas Bayes 250 years ago (Bayes and Price 1763); and the approach allows for flexibility in the strength assigned to the 20th century experience in the assessment of the 21st century Phillips curve. Heuristically, this strength could be termed the degree of conviction in the prior information. In the limiting case of essentially zero conviction in such prior experience, the Bayesian approach is equivalent to the approach in which a break is assumed in the Phillips curve in the 21st century and only post-1999 data is used to estimate the inflation process (i.e., the first approach above). Finally, the approach yields insights regarding the

information in macroeconomic data that are likely broadly relevant in macroeconomics—and therefore can inform future research on other questions. For example, the findings highlight the importance of assessing the strength of information in the data for assessing the evolution of other macroeconomic concepts, building on similar results in Kiley (2020a, 2020b).

The results that emerge from the analysis are very clear. Data over the period 2000–19 provide very little information with which to update prior views on the inflation process, at least with respect to inflation persistence or the “anchoring” of inflation. As a result, a Bayesian decisionmaker aware of the inflation process from 1960 through 1999—that is, endowed with a prior consistent with that process—would view the current inflation process as similar to that from 1960 through 1999 unless they place very little weight on the prior information. A direct implication of this result is that the inflation process may signal substantial persistence, suggesting the high inflation of 2021 may continue in 2022. The intuition for this finding is straightforward. Inflation was very stable from 2000 through 2019, which means the data witnessed few substantive deviations from its average or lagged values. Because the data contain few sizable deviations of inflation from its average or lagged values, the data provide little information regarding what would happen if inflation were to deviate sizably from its average or lagged values—an intuitive insight that also follows directly from the mathematics of a Bayesian least-squares regression. The lack of information in the data from 2000 through 2019 contrasts sharply with the precision of the prior view of the role of lagged inflation in the inflation process that is consistent with experience from 1960 through 1999, when inflation saw sizable swings away from its average value and experience suggests a sizable role for lagged inflation in the Phillips curve. This combination—low information in 2000–19 data and an informative prior consistent with 1960–99 experience—implies that the empirical analysis results in substantial inflation persistence unless the prior experience receives very little weight in the Bayesian decisionmaker’s calculus. These findings highlight how researchers may find it valuable to assess the information in their recent data using Bayesian methods in cases where there is prior information, as an approach to complement approaches such as time-varying parameter models or structural-break analyses.

Previous Literature. The analysis is related to the empirical literature examining the factors that determine the inflation process and how the importance of such factors may have shifted over time to explain high and variable inflation in the 1970s and low and stable inflation in the 21st century.¹ These include the degree to which the inflation process is “anchored” (i.e., the degree to which lagged inflation influences current inflation), the effect of unemployment on inflation (i.e., the slope of the Phillips curve), and the path of “supply shocks” or other supply factors that shift the relationship between inflation and economic slack. While alternative approaches are possible, a common—and simple—approach is to use a reduced-form Phillips curve relating inflation to its lags and the unemployment rate (e.g., Ball and Mazumder 2011, 2019; Gordon 2013; Kiley 2015b; or Blanchard 2016).

In this taxonomy, an “anchored” inflation process shows little effect of current inflation on subsequent inflation. The empirical work herein links anchoring to the persistence in inflation, which could reflect a variety of structural factors. For example, a sizable body of research has suggested that inflation persistence may have fallen in the 21st century (e.g., Williams 2006a, 2006b; Kiley 2008b, 2015b; Ball and Mazumder 2011, 2019; Stock 2011; Watson 2014; Coibion and Gorodnichenko 2015; Blanchard 2016; Carvalho et al. 2021; and Jorgenson and Lansing 2023). A prominent line of thought in this research is that inflation expectations (in the Phillips curve) followed an accelerationist structure in the decades before 2000—responding strongly to recent inflation experience—and that inflation expectations were anchored in the years after 2000—responding little to lagged inflation. But much of this work is similar to the approach herein, focusing on inflation persistence with little direct attention to expectations. Future work can consider the implications of the approach herein for expectations per se. Some research also has questioned the decline in inflation persistence (e.g., Pivetta and Reis 2007—although this study predates the low and stable inflation of the first two decades of the 21st century).

Research has also suggested a weaker relationship between unemployment and inflation in recent decades—that the Phillips curve

¹Theoretical modeling has also considered possible factors, e.g., Kiley (2007).

has flattened in the 21st century or earlier (Atkeson and Ohanian 2001). Ball and Mazumder (2011) suggest that “menu cost” models of nominal price and wage rigidity imply that such rigidities increase as the average rate of inflation falls, implying that more of the adjustment in nominal aggregate demand falls on output and less on inflation when inflation is low; this is exactly the finding emphasized in Kiley (2000), which analyzed support for this prediction across a large sample of countries. Research exploring the effects of downward nominal-wage rigidity points to a reduced effect of labor-market weakness on inflation in a low-inflation environment (Daly and Hobijn 2014). Kiley (2008b) and Boivin, Kiley, and Mishkin (2010) present evidence that a clear commitment to price stability in recent decades, in the form of a monetary policy rule with a more sizable response to inflation, acts to substantially stabilize inflation expectations and mitigate fluctuations in inflation. Such a shift in monetary policy behavior is consistent with an observed flattening in the Phillips curve in the 21st century. Del Negro et al. (2020) find a similar role for monetary policy but find a large role for structural factors related to aggregate supply in the flattening of the Phillips curve.

The literature pursues different empirical approaches, within or outside a Phillips curve approach. Much of the literature considers reduced-form Phillips curves estimated across subsamples of the data—i.e., considers breaks in estimated equations. Examples include Williams (2006b), Kiley (2008b, 2015b), Blanchard (2016), and Ball and Mazumder (2019). Other work explicitly models time variation in parameters, for example, in vector autoregressions as in Cogley and Sargent (2005) and Primiceri (2005). A particularly influential class of time-varying parameter models are time-varying parameter unobserved component models (Stock and Watson 2007; and Kiley 2008a). In this approach, it is common for results to suggest that the variance of the permanent drift components was lower in the 2000–19 period than earlier—i.e., that inflation was anchored in the 21st century.

The approach herein uses the textbook approach to Bayesian regression (e.g., Kim and Nelson 1999, Chapter 7) to examine the information content of the data for parameters of a Phillips curve relative to the information in a reasonable prior. This approach has not been used in discussions of inflation dynamics. Bayesian approaches

are often used in estimation of time-varying parameter models, as in Primiceri (2005)—but these analyses do not focus on the information content in the data relative to that in the prior; rather, they emphasize the value of Bayesian methods given the complexity of estimating such models. Kiley (2020a, 2020b) highlights how macroeconomic relationships may be poorly informed by aggregate time-series data and how the available data may not lead posterior assessments to differ from prior views in an examination of the equilibrium real interest rate.²

The analysis also raises questions regarding why inflation was more stable from 2000 onward. The Bayesian approach indicates that the data do not contain information to suggest a very large change in inflation persistence relative to pre-2000 experience. Under this view, the post-2000 inflation experience would be ascribed to “luck” that resulted in smaller shocks to inflation. Previous research has noted the challenges associated with distinguishing “luck” from structural changes (e.g., Ahmed, Levin, and Wilson 2004).

Structure of the Remaining Sections. Section 2 discusses data and the framework a Bayesian decisionmaker endowed with a prior for the inflation process in the United States consistent with data over the second half of the 20th century would use to incorporate the information from the data over the 2000–19 period in their view on the inflation process. Section 3 presents results, intuition, and implications. Section 4 concludes.

2. Data and Approach

2.1 Data

The study analyzes inflation in the United States. The analysis focuses on the consumer price index (CPI), produced monthly by the U.S. Bureau of Labor Statistics. The focus of the investigation is the evolution of the persistence of inflation and, to a lesser extent, the slope of the Phillips curve. To abstract from the

²Given limited information in aggregate time-series data, Fitzgerald and Nicolini (2014) and Kiley (2015a) analyze Phillips curves using city-level and state-level data, respectively. Hooper, Mishkin, and Sufi (2020) and Hazell et al. (2022) build on this approach.


volatility induced by fluctuations in food and energy prices, the empirical work uses the CPI excluding food and energy (core CPI), which is available from January 1957 to December 2021. The results are generally similar for the overall CPI, reflecting the correlation between overall and core CPI (e.g., Kiley 2008a). The results are also similar when the price index considered is the chain-weighted price index for personal consumption expenditures (PCE prices). Note that the Federal Open Market Committee (FOMC) of the Federal Reserve System has defined its inflation objective of 2 percent in terms of PCE prices since 2012 and inflation as measured by the CPI index has averaged a few tenths above inflation as measured by PCE prices in recent decades; for this reason, we will refer to the inflation objective in the United States, as measured by the CPI, as in the range of 2 percent.

The Phillips curve framework relates inflation to a measure of economic slack. The analysis uses the unemployment rate of the civilian non-institutional population aged 16 and over (the unemployment rate), produced monthly by the U.S. Bureau of Labor Statistics.

Figure 1 presents the data on inflation and the unemployment rate. The inflation measure presented is the 12-month change in the natural logarithm of the core CPI (top panel). Inflation was low and stable in the late 1950s and early 1960s. Inflation rose over the late 1960s and was both higher and more volatile over the 1970s. After the tightening in monetary policy associated with the Volcker disinflation that began in late 1979, inflation drifted lower over the course of the 1980s and early 1990s. Over this period from the late 1950s through the mid-1990s, inflation appeared to be persistent—that is, years in which inflation exceeded the average over this period tended to be followed by years in which inflation was above its average. From 2000 until 2019, inflation was generally low—near 2 percent—and stable. Inflation jumped out of its 2000–19 range in 2021, reaching about 5½ percent in a 12-month basis in December 2021. The unemployment rate (bottom panel) rises sharply during recessions and declines during expansions, highlighting how it is a good measure of the state of the U.S. business cycle.

Table 1 presents some summary statistics on inflation and the unemployment rate. Statistics are presented for monthly data and

Figure 1. CPI Inflation and the Civilian Unemployment Rate



Source: Bureau of Labor Statistics and author’s calculations.

for data on an annual average basis, as the Phillips curve analysis will consider data at both the monthly and annual frequency as one robustness check. Statistics are shown for three sample periods: late 1950s–2019, late 1950s–1999, and 2000–19. These three sample periods will be referred to as the full sample, the pre-2000 sample, and the post-1999 sample. The years 2020 and 2021 are excluded from the table and the estimation sample, reflecting the unprecedented (and unusual) effects of the COVID-19 pandemic; econometric work will almost surely explore various ways to treat these unusual years in emerging research.³ Two aspects of the summary statistics will prove important in understanding the results. First, inflation was much more volatile in the pre-2000 period than in the post-1999

³Lenza and Primiceri (2020) and Schorfheide and Song (2021) highlight the potential sensitivity of macroeconomic estimates to the COVID-19 pandemic period, with the former suggesting some approaches to handling these issues.

Table 1. Data Summary Statistics

	Observations	Mean (Annual Rate)	Std. Deviation	Autocorrelation
Monthly Data				
<i>Full Sample</i>				
CPI Inflation (Percent)	744	3.6	0.25	0.66
Unemployment Rate (Percent)		6.0	1.6	0.99
<i>Pre-2000 Sample</i>				
CPI Inflation (Percent)	504	4.3	0.27	0.62
Unemployment Rate (Percent)		6.0	1.5	0.99
<i>Post-1999 Sample</i>				
CPI Inflation (Percent)	240	2.0	0.08	0.21
Unemployment Rate (Percent)		5.9	1.8	0.99
Annual Data				
<i>Full Sample</i>				
CPI Inflation (Percent)	61	3.6	2.4	0.88
Unemployment Rate (Percent)		6.0	1.6	0.80
<i>Pre-2000 Sample</i>				
CPI Inflation (Percent)	41	4.4	2.6	0.83
Unemployment Rate (Percent)		6.0	1.5	0.78
<i>Post-1999 Sample</i>				
CPI Inflation (Percent)	20	2.0	0.4	0.42
Unemployment Rate (Percent)		5.9	1.8	0.77
Source: Bureau of Labor Statistics and author's calculations.				
Note: Full sample: Jan. 1958–Dec. 2019 (monthly) and 1959–2019 (annual); Pre-2000 sample: Jan. 1958–Dec. 1999 (monthly) and 1959–99 (annual); Post-1999 sample: Jan. 2000–Dec. 2019 (monthly) and 2000–19 (annual).				

period, as can be seen in the standard deviations of the series during the periods. Second, inflation was much more persistent in the pre-2000 period than in the post-1999 period, as can be seen in the autocorrelation of the series.

2.2 Empirical Approach

The analysis considers estimates of a Phillips curve in which inflation ($\Delta p(t)$) depends on its own lags and the (lagged) unemployment rate ($u(t)$) as in Equation (1) (in which a constant term is suppressed):

$$\Delta p(t) = \sum_{j=1}^{N-1} b(j) \cdot \Delta p(t-j) + a \cdot u(t-1) + e(t). \quad (1)$$

$b(j)$ are the coefficients governing persistence, a is the slope of the Phillips curve, and $e(t)$ is the residual reflecting “supply” shocks and other unmodeled factors which follows a normal distribution ($e(t) \sim N(0, \sigma^2)$). An anchored Phillips curve would tend to have small $b(j)$, whereas an unanchored Phillips curve will tend to have large $b(j)$. The sum of these coefficients, $\sum_{j=1}^N b(j)$, will be the statistic of focus, with a sum near 1 representing an unanchored accelerationist Phillips curve.

The vector of coefficients in Equation (1) is denoted by Γ . The matrix containing the dependent variable (inflation) will be denoted $Y(T \times 1)$, where T is the number of observations, the matrix of right-hand-side variables will be denoted $X(T \times N)$, and the matrix of error terms will be denoted $E(T \times 1)$, yielding

$$Y = X\Gamma + E. \quad (2)$$

The classical approach to inference would estimate Γ by least squares as $\Gamma^{LS} = (X'X)^{-1}X'Y$.

The decisionmaker herein follows a Bayesian approach. The decisionmaker is endowed with a prior for Γ that is given by the normal distribution with mean $\tilde{\Gamma}$ and variance-covariance matrix V —i.e., a prior distribution $\Gamma \sim N(\tilde{\Gamma}, V)$. The analysis proceeds under the assumption that the decisionmaker knows the variance of $e(t)(\sigma^2)$, so the Bayesian estimation is conditional on σ^2 and the

assumed prior is the natural conjugate prior for Γ conditional on σ^2 . An alternative approach would also specify a prior view on σ^2 and jointly estimate the posterior distributions of Γ and σ^2 . This alternative approach yields essentially identical results for reasonable priors on σ^2 but yields more complicated algebraic expressions that slightly impede intuition for those less familiar with Bayesian least squares. As a result, the simpler approach is adopted herein.

Given the prior information, the decisionmaker estimates Γ by combining their prior information and the data—i.e., the prior distribution and the likelihood function of the data—to form the posterior distribution for Γ and estimates Γ to maximize this posterior distribution. This is a textbook example of Bayesian regression (Kim and Nelson 1999, Chapter 7), with the resulting estimate $\hat{\Gamma}$ given by

$$\hat{\Gamma} = (V^{-1} + \sigma^{-2} X'X)^{-1} (V^{-1} \tilde{\Gamma} + \sigma^{-2} X'X \Gamma^{LS}). \quad (3)$$

Notice in Equation (3) that the Bayesian decisionmaker estimates the parameters as the matrix-weighted average of their prior information and the least-squares estimate, with weights given by the precision of the information in the prior and the data (e.g., by the inverses of the variance-covariance matrices of $\tilde{\Gamma}$ and Γ^{LS} , V and $\sigma^2(X'X)^{-1}$).

Equation (3) suggests a natural approach to considering different degrees of conviction regarding the prior information. The Bayesian decisionmaker can further “weight” their prior by a factor w , as in

$$\hat{\Gamma} = (wV^{-1} + (1-w)\sigma^{-2} X'X)^{-1} (wV^{-1} \tilde{\Gamma} + (1-w)\sigma^{-2} X'X \Gamma^{LS}). \quad (4)$$

Intuitively, essentially zero weight on the prior information returns the least-squares estimate. This “weighting” terminology is convenient. Mathematically, it is equivalent to considering a less informative prior. Specifically, estimates with a weight of w on the prior information are equivalent to estimates with a prior for Γ with the same mean and a variance-covariance matrix equal to $\frac{1-w}{w}V$. For example, a weight on the prior equal to 20 percent ($w = 0.2$) is equivalent to a prior with a four-times looser variance-covariance matrix $4V$.


With this background, the approach involves choosing the lag specification in the Phillips curve, a choice of the prior distribution and consideration of alternative weights on this prior distribution.

- *Lag Specification in the Phillips Curve:* In the estimates using monthly data, the lag length equals 12 ($N = 13$) and the coefficients on the 1st through 12th lag are equal—i.e., $\sum_{j=1}^{N-1} b(j) \cdot \Delta p(t-j) = b(1) \sum_{j=1}^{12} \Delta p(t-j)/12$; alternative choices for the lag structure yielded similar results (see Section 3.3), and this specification is simplest. For the estimates using annual data, the lag length equals 1 ($N = 2$).
- *Choice for Prior $\Gamma \sim N(\tilde{\Gamma}, V)$:* The prior distribution used to inform estimates of the 21st century Phillips curve is given by the values consistent with the pre-2000 sample. $\tilde{\Gamma}$ is given by the least-squares estimate for this sample and V is the associated variance-covariance matrix. This is akin to an empirical Bayesian approach. The thought experiment is one in which the decisionmaker was endowed with information on the Phillips curve in the latter half of the 20th century and chooses to update their view following the realization of data from 2000 through 2019.
- *Choice of Weights w :* To consider decreasing levels of conviction in the relevance of the 20th century prior (i.e., looser priors), four values for weights on the prior are considered, with the factor w taking values of 0.5, 0.2, 0.05, or (approximately) 0—corresponding to variance covariance matrices for the prior equal to V , $4V$, $19V$, and an uninformative prior.

Figure 2 presents the prior distributions for the coefficient on the lags of inflation and the slope of the Phillips curve for these alternative weights. The priors show high values of persistence (a central tendency for the sum on inflation lags near 1) and a notable slope to the Phillips curve (i.e., a negative slope with a degree of precision in the prior distribution).⁴

⁴The estimates use inflation data expressed at annual rates. This convention makes the slope coefficients somewhat more comparable (albeit still not strictly comparable, reflecting how time averaging would affect impact estimates at different frequencies).

Figure 2. Prior Information



Source: Bureau of Labor Statistics and author’s calculations.


3. Results

3.1 Estimates

The results for posterior estimates of the Phillips curve parameters conditional on data from 2000 through 2019 by the Bayesian decisionmaker are shown in Figure 3 and Table 2. Two results emerge clearly.

First, the degree of persistence in the posterior estimates is high in all cases except those with very little weight on the 20th century prior. In the monthly estimates, the coefficient on lagged inflation exceeds 0.9 with equal weights on prior and data, exceeds 0.85 when the weight on the prior is 0.20, and is near 0.7 when the weight on the prior is 0.05; in the limiting case of essentially no weight on the prior, the coefficient on the lags is small at about 0.2 (in line with

Figure 3. Posterior Estimates



Source: Bureau of Labor Statistics and author’s calculations.

simple least squares for the 2000–19 period). For the annual data, the coefficient on lagged inflation in the posterior is about $7/8$, $3/4$, and $1/2$ for weights on the prior of 0.5, 0.2, and 0.05—whereas the coefficient is 0 for the uninformative prior.

Second, the slope of the Phillips curve is consistently smaller in absolute value in the posterior estimates of the parameters—irrespective of the weight on the prior. For example, the slope of the Phillips curve is less than $1/2$ the prior value (the pre-2000 data value from least squares) in the posterior estimates for all weights on the prior.

These results suggest that the approach of a Bayesian decision-maker confirms the finding in the literature that the Phillips curve is “flatter.” In contrast, the approach of a Bayesian decisionmaker does not find the degree of “anchored” inflation as would be implied by an estimate using only recent data, in the sense that the coefficient on lagged inflation is substantial when prior information is

Table 2. Estimation of Posterior of Parameters

$$\Delta p(t) = b(1) \sum_{j=1}^N \Delta p(t-j)/N + a \cdot u(t-1) + e(t)$$

incorporated. Nonetheless, the coefficient on the lags of inflation is less than 1—so the process is not of the “accelerationist” type and would generally imply that deviations of inflation from its average may be long-lived but are ultimately transitory.

3.2 *Intuition*


The results are clear—but they are also very intuitive. Recall from Equation (4) that the posterior estimates are the matrix-weighted average of the prior mean and least-squares estimate for the post-1999 period, with weights given by the inverse of the variance-covariance matrices.

For the slope of the Phillips curve, the variance-covariance matrix of the post-1999 least-squares estimate implies a fair degree of precision. Recall that this matrix is $\sigma^2(X'X)^{-1}$ and its inverse is $\sigma^{-2}X'X$. The component of this matrix that “weights” the least-squares estimate is dominated by the sum of squared deviations of the unemployment rate from its mean (assuming a modest covariance between inflation and unemployment). The summary statistics in Table 1 show that this sum of squares remains sizable relative to its pre-2000 value, as indicated by the standard deviation of the unemployment rate. As a result, the least-squares estimate receives considerable weight in the Bayesian decisionmaker’s calculus.

In contrast, inflation is quite stable in the post-1999 period. Its sum of squared deviations from the mean is modest relative to the pre-2000 experience, as indicated by the standard deviation. This implies the data for 2000–19 receive relatively little weight in a Bayesian decisionmaker’s calculus when assessing the persistence of inflation. In words, inflation did not deviate from its average value much over 2000–19, and hence a Bayesian decisionmaker does not weigh experience over that period highly when evaluating how persistent a deviation of inflation from its mean is likely to be. This is intuitive—the data do not provide examples of what would happen should inflation deviate from its mean, and hence the data are not informative about what would happen following such a deviation.

A look at measures of fit provides some intuition for why models with such different dynamic properties—the estimates with different weights on the prior—emerge. Figure 4 presents a dynamic simulation of the estimates from 2000 through 2019 in the top panel and

Figure 4. Visual Assessments of Fit



Source: Bureau of Labor Statistics and author’s calculations.

the residuals implied by each estimate in the bottom panel, in both cases for the monthly-data specification. All the estimates lead to a dynamic simulation of relatively low and stable inflation—although the case with a weight of 0.5 on the prior information shows more variability and a worse “fit” than the others. Looking at one-period misses, the residuals are all very highly correlated. The residuals are highly correlated across specifications with a large coefficient on the lag and a small coefficient on the lag because inflation has been stable and the contribution of the lag is small, irrespective of the coefficient. This intuition is the same as that above—inflation has been stable and hence the data do not differentiate much between a specification with a large or small coefficient on lagged inflation.

This intuition also provides insight into the comparison of the full sample results from least squares (reported in Table 2) with the results from a Bayesian approach. The full-sample results are similar to those of a Bayesian decisionmaker that places weight on the pre-2000 experience—inflation is persistent, and the Phillips curve is flatter than in the pre-2000 period. The pre-2000 experience

dominates the variation in the data and hence drives the full-sample estimates of persistence. This comparison also highlights a potential weakness of the approach of a Bayesian decisionmaker for a researcher that views a structural break as likely. The Bayesian decisionmaker views the parameters as drawn from a stable distribution and allows the data to move them away from their prior view. This approach is consistent with the notion that the future may look different than embedded in prior information, but not consistent with a structural break—which would imply that prior information has no value. The findings herein suggest that a researcher may wish to entertain the possibility that experience from 2000 through 2019 is consistent with a sizable degree of persistence in inflation, but also may wish to consider alternative approaches that allow for structural breaks or time-varying parameters.

3.3 *Robustness and Implications*

As noted above, the basic results do not depend upon the specific lag structure assumed for inflation in the results reported in Table 2. To illustrate the robustness of the results to alternative specifications, Table 3 considers a slightly more flexible lag structure, as in

$$\Delta p(t) = b(1) \sum_{j=1}^3 \frac{\Delta p(t-j)}{3} + b(2) \sum_{j=4}^{12} \frac{\Delta p(t-j)}{9} + a \cdot u(t-1) + e(t).$$

In this alternative, the sum of the coefficients on the lags of inflation ($b(1) + b(2)$) gives a rough gauge of the persistence of inflation. The results are substantially similar to those for the simpler specification in Table 2, with $b(1) + b(2)$ estimated at essentially the same values as those for $b(1)$ in Table 2.

The implications of the results for inflation forecasts are direct. A Phillips curve based on post-1999 data alone would imply a sharp deceleration of inflation in 2022, as there is very little persistence estimated in that case. In contrast, the Bayesian estimates imply that inflation will remain quite high in 2022 in the absence of unexpected shocks. Generally speaking, the results herein suggest a higher degree of persistence is plausible, pointing to potentially higher inflation in 2022.

4. Conclusions

This analysis examined the extent to which the experience from 2000 to 2019 should lead a Bayesian decisionmaker to update their assessment of inflation dynamics. Given a prior for inflation dynamics consistent with 1960–99 data, a Bayesian decisionmaker would not update their view of inflation persistence in light of 2000–19 data unless they placed very low weight on their prior information. In other words, 21st century data contain very little information to dissuade a Bayesian decisionmaker of the view that inflation fluctuations are persistent, or “unanchored.”

The idea that data over short sample periods may provide limited information regarding macroeconomic relationships may be relevant for other areas in macroeconomics. For example, Kiley (2020a, 2020b) finds modest information in the data for estimates of the equilibrium real interest rate. These findings suggest macroeconomists may find useful a Bayesian approach that examines the information in the data more thoroughly than is common in empirical macroeconomics.

References

- Ahmed, S., A. Levin, and B. A. Wilson. 2004. “Recent U.S. Macroeconomic Stability: Good Policies, Good Practices, or Good Luck?” *Review of Economics and Statistics* 86 (3): 824–32. <https://EconPapers.repec.org/RePEc:tpr:restat:v:86:y:2004:i:3:p:824-832>.
- Atkeson, A., and L. E. Ohanian. 2001. “Are Phillips Curves Useful for Forecasting Inflation?” *Quarterly Review* (Federal Reserve of Minneapolis) 25 (1): 2–11.
- Ball, L., and S. Mazumder. 2011. “Inflation Dynamics and the Great Recession.” *Brookings Papers on Economic Activity* (Spring): 337–81.
- . 2019. “A Phillips Curve with Anchored Expectations and Short-Term Unemployment.” *Journal of Money, Credit and Banking* 51 (1): 111–37.
- Bayes, T., and R. Price. 1763. “An Essay Towards Solving a Problem in the Doctrine of Chance. By the late Rev. Mr. Bayes, communicated by Mr. Price, in a letter to John Canton, A. M. F. R. S.”

- Philosophical Transactions of the Royal Society of London* 53: 370–418. doi:10.1098/rstl.1763.0053.
- Blanchard, O. 2016. “The Phillips Curve: Back to the ’60s?” *American Economic Review* 106 (5): 31–34.
- Boivin, J., M. T. Kiley, and F. S. Mishkin. 2010. “How Has the Monetary Transmission Mechanism Evolved Over Time?” In *Handbook of Monetary Economics*, ed. B. M. Friedman and M. Woodford, Edition 1, Vol. 3, 369–422 (chapter 8). Elsevier.
- Carvalho, C., S. Eusepi, E. Moench, and B. Preston. 2021. “Anchored Inflation Expectations.” Available at SSRN 3018198.
- Cogley, T., and T. J. Sargent. 2005. “Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.” *Review of Economic Dynamics* 8 (2, April): 262–302.
- Coibion, O., and Y. Gorodnichenko. 2015. “Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation.” *American Economic Journal: Macroeconomics* 7 (1, January): 197–232.
- Daly, M. C., and B. Hobijn. 2014. “Downward Nominal Wage Rigidities Bend the Phillips Curve.” Working Paper No. 2013-08, Federal Reserve Bank of San Francisco.
- Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti. 2020. “What’s Up with the Phillips Curve?” NBER Working Paper No. 27003.
- Fitzgerald, T. J., and J. P. Nicolini. 2014. “Is There a Stable Relationship between Unemployment and Future Inflation? Evidence from U.S. Cities.” Working Paper No. 713, Federal Reserve Bank of Minneapolis.
- Gordon, R. J. 2013. “The Phillips Curve is Alive and Well: Inflation and the NAIRU During the Slow Recovery.” NBER Working Paper No. 19390, August.
- Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson. 2022. “The Slope of the Phillips Curve: Evidence from U.S. States.” *Quarterly Journal of Economics* 137 (3): 1299–1344.
- Hooper, P., F. S. Mishkin, and A. Sufi. 2020. “Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or Is It Just Hibernating?” *Research in Economics* 74 (1): 26–62.
- Jorgensen, P., and K. J. Lansing. 2023. “Anchored Inflation Expectations and the Flatter Phillips Curve.” Working Paper No. 2019-27, Federal Reserve Bank of San Francisco.

- Kiley, M. T. 2000. "Endogenous Price Stickiness and Business Cycle Persistence." *Journal of Money, Credit and Banking* 32 (1, February): 28–53.
- . 2007. "Is Moderate-to-High Inflation Inherently Unstable?" *International Journal of Central Banking* 3 (2, June): 173–201.
- . 2008a. "Estimating the Common Trend Rate of Inflation for Consumer Prices and Consumer Prices Excluding Food and Energy Prices." Finance and Economics Discussion Series No. 2008-38, Board of Governors of the Federal Reserve System.
- . 2008b. "Inflation Expectations, Uncertainty, the Phillips Curve, and Monetary Policy — Comments." In *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective*, ed. J. Fuhrer, J. S. Little, G. P. Olivei, and Y. K. Kodrzycki. Cambridge, MA: MIT Press.
- . 2015a. "An Evaluation of the Inflationary Pressure Associated with Short- and Long-Term Unemployment." *Economics Letters* 137 (December): 5–9. ISSN 0165-1765. <http://dx.doi.org/10.1016/j.econlet.2015.10.005>.
- . 2015b. "Low Inflation in the United States: A Summary of Recent Research." FEDS Notes, Board of Governors of the Federal Reserve System, November 23. <https://doi.org/10.17016/2380-7172.1665>.
- . 2020a. "The Global Equilibrium Real Interest Rate: Concepts, Estimates, and Challenges." *Annual Review of Financial Economics* 12 (1, December): 305–26.
- . 2020b. "What Can the Data Tell Us about the Equilibrium Real Interest Rate?" *International Journal of Central Banking* 16 (3, June): 181–209.
- Kim, C. J., and C. R. Nelson. 1999. *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*. Cambridge, MA: MIT Press.
- Lenza, M., and G. E. Primiceri. 2020. "How to Estimate a VAR after March 2020." NBER Working Paper No. 27771.
- Pivetta, F., and R. Reis. 2007. "The Persistence of Inflation in the United States." *Journal of Economic Dynamics and Control* 31 (4): 1326–58. <https://EconPapers.repec.org/RePEc:eee:dyncon:v:31:y:2007:i:4:p:1326-1358>.

- Primiceri, G. E. 2005. "Time Varying Structural Vector Autoregressions and Monetary Policy." *Review of Economic Studies* 72 (3): 821–52.
- Schorfheide, F., and D. Song. 2021. "Real-Time Forecasting with a (Standard) Mixed-Frequency VAR during a Pandemic." NBER Working Paper No. 29535. <http://www.nber.org/papers/w29535>.
- Stock, J. 2011. "Comment on Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity* (Spring): 387–402.
- Stock, J. H., and M. W. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39 (s1, February): 3–33.
- Watson, M. W. 2014. "Inflation Persistence, the NAIRU, and the Great Recession." *American Economic Review* 105 (5, May): 31–36.
- Williams, J. C. 2006a. "Inflation Persistence in an Era of Well-Anchored Inflation Expectations." FRBSF Economic Letter No. 2006-27, Federal Reserve Bank of San Francisco.
- . 2006b. "The Phillips Curve in an Era of Well-Anchored Inflation Expectations." Unpublished Working Paper, Federal Reserve Bank of San Francisco.

Researching the Research: A Central Banking Edition*

Simona Malovaná,^{a,b} Martin Hodula,^{a,c} and Zuzana Gric^{a,d}

^aCzech National Bank

^bPrague University of Economics and Business

^cVSB—Technical University of Ostrava

^dMasaryk University in Brno

We build two unique data sets on the research in central banks in Europe and the United States, offering a novel insight into central banks' research activities, the topics covered, collaborations networks, gender diversity, and research popularization. Even though we identify significant heterogeneity among different regions, we reveal several important stylized facts. Following the global financial crisis, financial stability surpassed monetary policy as the leading research topic. We have also documented a substantial decline in papers with single authors, from 40 percent in 2000 to less than 20 percent in 2019. Still, research in central banks is highly concentrated, as the top 10 percent of authors contribute to about 50 percent of all central banks' research publications. While central banks form enormous research networks, we find that most of this research collaboration is region specific with an increasing representation of women. Nevertheless, the gender gap persists and is closing only slowly. In this respect, small central banks are found to employ more female researchers than large ones. Last but not least, major central banks with a well-established research tradition achieve the highest average impact factor, with a few research papers contributing the most to this average.

JEL Codes: A1, A3, D85, E58, O31.

*The authors note that the paper represents their own views and not necessarily those of the Czech National Bank. We would like to thank Dennis Essers, Jan Brůha, two anonymous referees, and participants of the Czech National Bank's seminar for useful comments. All errors and omissions remain the fault of the authors. Part of the research behind this paper was funded by the Prague University of Economics and Business institutional research support grant no. IP100040 and Technical University of Ostrava grant no. SP2022/38. Simona Malovaná (corresponding author): simona.malovana@cnb.cz; Martin Hodula: martin.hodula@cnb.cz; Zuzana Gric: zuzana.gric@cnb.cz.

1. Introduction

Economic and financial research plays a pivotal role in central banks around the world. Research units are tasked with providing policy-makers with inputs which help to expand the knowledge base needed for the central bank's core activities. Given the fact that central banking tasks continue to grow in complexity, the importance of research and the demands placed on it are expected to grow too. Thus, the question is not whether to conduct research in central banks, but how to organize it, how much to invest in it, and how to assure effective transmission of research outcomes into the policy decisionmaking process. Nowadays, an increasing number of central banks have their own research divisions or whole departments staffed with both junior- and senior-level research economists.¹ This is a manifestation of the general belief that sound policy must build on cutting-edge economic thinking. Moreover, research in central banks has established itself as a prominent contributor to the academic literature over time. Claveau and Dion (2018) show that major monetary economics journals are increasingly dominated by central banks' research papers. Another interesting trend is the visibly increasing share of non-U.S. central bankers who have been invited to the annual Jackson Hole Economic Policy Symposium.² This has allowed them to conduct high-level discussions about current topics, to evaluate the ex ante appropriateness of intended measures, and to analyze the efficiency of measures applied.

The literature on central banks (and, more generally, on independent expert organizations) supplies at least two hypotheses that might explain the number of resources going to research today in central banks. First, being perceived as "scientific" organizations might be a huge asset for central banks in their interactions with external actors (elected officials, other regulatory bodies, market

¹The high level of specialization and proficiency in economics and finance is also apparent from the governor appointments. Lebaron and Dogan (2016) analyze the biographies of 312 incumbent and former central bank governors from across the world, observing a high qualification level, predominantly in economics.

²The proportion of market participants fell from 27 percent in 1982 to 3 percent in 2013 as they were "replaced" by non-U.S. central bankers, whose share increased from 3 percent to 31 percent (*The Economist* 2014).

participants, and the general public). High-quality research publications can enhance a central bank's reputation and increase its credibility, making it easier to defend policy actions both publicly and privately.³ Second, there is an international dynamics inside the central banking community that pushes toward more research through mimicry.⁴ The observed scientization of central banks might be akin to a bandwagon effect: a few research-oriented central banks are identified as "successes," which incentivizes other central banks to copy the model and, as the model diffuses, it becomes the thing to do, rationalized by the claim that in-house research makes for better policy (Cecchetti 2002; Trichet 2007; Bolton, Cecchetti, and Reichlin 2014; Carney 2014).⁵

In our paper, we provide an international overview of the research conducted in central banks. We perform our exploration across both the cross-sectional and the time dimension. Specifically, we aim to shed some light on how central banks organize and communicate their research activities, how they perform in terms of research outputs, and what the recent trends are in areas such as the research topics covered, collaborations between central banks and with other institutions, and gender diversity. For this purpose, we create two unique data sets. The first one contains information on central banks' research activities in three main areas: research publications, popularization of research, and organization of research. Each area is evaluated on a number of different factors, together making a list of 26 items. The second data set comprises information about more than 20,000 research papers published in European and U.S. central banks over the period of 2000–19. To the best of our knowledge, neither of these two data sets has previously been constructed. As such, they offer a unique insight into central banks' research activities, allowing us to explore heterogeneity

³Literature shows that institutions craft their rhetoric to increase their popularity and protect themselves from audience-based pressures (Maor, Gilad, and Bloom 2013; Gilad, Maor, and Bloom 2015; Moschella and Pinto 2019).

⁴For example, in the early stage of setting up the European System of Central Banks, the U.S. Federal Reserve served as the primary template for research organization (Fase and Vanthoor 2000).

⁵Similar speeches on the importance of research in central banks were given at the end of the millennium as well—see Taylor (1998).

among central banks from different regions and changes in research conducted before and after the global financial crisis (GFC) of 2007–09. Documenting these issues may help central banks to identify best practices.

We complement the existing body of literature along at least three dimensions. First, we add to the literature examining the research conducted in academia. This includes studies that analyze the formation of research networks among academics (Adams 2012) and studies commenting on the demographics and trends in the articles published in economic journals (Card and DellaVigna 2013; Hamermesh 2013). Our paper verifies whether some of the trends observed in academic research regarding the number of coauthors, network density, distribution of authorship, and others are also present in central banks' research. In this respect, our findings largely complement those of Essers, Grigoli, and Pugacheva (2020), who employed data on 6,152 International Monetary Fund (IMF) working papers written in the period of 1990–2017, including authors' names and e-mails, number of pages and citations, and JEL codes, using the IMF website catalogue. In addition, they complemented the data with employee-level information. Our analysis offers a rather different perspective, as it includes an international sample of central banks and supranational institutions, while it supports many of the findings of Essers, Grigoli, and Pugacheva (2020). Second, our paper echoes the literature on central bank communication (Dincer and Eichengreen 2007; Blinder et al. 2008) by offering a unique view on how research—one of the underlying processes behind policy decisionmaking—is conducted. Third, we contribute to the growing body of literature on gender gaps in central banks (Charl  ty, Romelli, and Santacreu-Vasut 2017; Diouf and P  pin 2017) by studying the gender structure of research teams.

The rest of the paper is organized as follows. Section 2 describes two unique data sets on central banks' research. Section 3 discusses the main purpose of research in central banks and categorizes the main research activities performed by them. Section 4 gets more into the detail and presents a few stylized facts on the topics researched in central banks, authorship networks, the gender structure of research teams, and the impact factor of central banks' publications. Section 5 concludes.

2. Data Sets on Central Banks' Research Activity

We construct two unique data sets allowing us to explore various attributes of research in central banks.⁶ We focus our attention on two regions—the United States and Europe. The U.S. Federal Reserve is a reference on its own in the world of central banks, especially in terms of its research (Claveau and Dion 2018). Numerous testimonies from European central banks (the European Central Bank notwithstanding) confirm that they aim(ed) to emulate the Fed and the Anglo-Saxon tradition in terms of research powers (Fase and Vanthoor 2000; Mudge and Vauchez 2016). The two regions thus form a good sample to study, as they are likely to share important similarities.⁷ According to the IDEAS/RePEc ranking of top 25 percent central banks in terms of research, our sample covers 28 out of 32 central banks listed as of January 2022. Given such a large coverage, the patterns found for our sample of central banks may be generalized also for other jurisdictions. We complement these regions with data on two additional international institutions, namely the Bank for International Settlements (BIS) and the European Central Bank (ECB, also including the European Systemic Risk Board, ESRB), as both encompass research activities serving policy conduct in central banks.⁸ Altogether, we compile information on 55 central banks. We include all 12 regional Federal Reserve Banks in the Federal Reserve System and the Federal Reserve Board (FRB) and 40 European national central banks. Within Europe, we further distinguish between three regional areas—the euro area original members (EA orig.), Western and Northwestern Europe (W & NWE), and Eastern and Southeastern Europe and the Baltic Countries (E & SEE, BC; see Table 1). We hypothesize that research quantity, quality,

⁶Both data sets, including code used for their construction, are available at <https://simonamalovana.org/>.

⁷In the construction of the data sets, we needed to find a compromise between data accessibility and a related amount of labor and relative sample homogeneity. Recent data suggest that economists from the U.S. Federal Reserve Banks, international financial institutions, and euro zone central banks are cited more frequently than economists with similar characteristics from central banks located in emerging markets (Rybacki and Serwa 2021).

⁸We deliberately restrict other high-ranked research-oriented institutions from our sample, such as the IMF or the World Bank, because they are not perceived as central banks nor are their functions related to central bank policy conduct.

and collaboration, among other things, will be different between the regions. For instance, we assume a more dense collaboration network within the euro area original members, reflecting a longer period of cooperation via various working groups and research networks. Furthermore, we assume the higher quality of research conducted by central banks in W & NWE relative to E & SEE, BC due to a long tradition of research in both central banks and academia. As such, the differentiation between these regions will help us to explore the potential heterogeneity.

The first data set leverages information available on each of the 55 central banks' websites and social media accounts. The data set (hereafter referred to as the Research Score Data Set) comprises information on central banks' research activities in three main areas: research organization, publications, and popularization. Each area is evaluated using several different factors, together making a list of 26 items (see Table A.1 in Appendix A). The items include, for example, the types of research publication series, the extent of research opportunities offered, the level of autonomy of the research departments, and the degree to which U.S. and European central banks use their websites, social media (Twitter, LinkedIn, and YouTube) and additional forms of communication such as blogs, newsletters, and bulletins to promote the research. Out of the 26 items, 21 were harvested from banks' websites and four from banks' social media accounts, and the last item was collected as a single observation from the IDEAS/RePEc website (see the next paragraph). The data were collected during the observation period of February–March 2020 and should therefore be interpreted as of these months. Each item is assigned a score of between 0 and 1: 1 if the bank engages in the given activity in full; 0.50/0.25 if the activity is in place only partially; and 0 otherwise. To obtain a separate measure for each of the three areas and a single overall measure for all 26 items, the research scores are constructed as the sum of the scores assigned to the individual items; the individual items are not weighted. As such, we get an *overall research score*, a *research organization score*, a *research publication score*, and a *research popularization score*. All the scores are calculated on the level of individual central banks and visualized in the form of a geographical heat map (see Figure 2 and Figure B.1 in Appendix B).

The second data set comprises information about more than 20,000 research papers published in 50 central banks over the period of 2000–19.⁹ This data set was extracted from the IDEAS/RePEc website using the web-scraping technique, and it includes the paper's title, the authors' names and reported affiliations, the abstract, keywords and JEL codes, and an indication of whether the paper has been published in a scientific journal.¹⁰ We harvested information on working paper series, discussion paper series, and occasional paper series, and on other research publication series similar to those three (for instance, research publication series with a special focus, such as Monetary Policy Committee (MPC) discussion papers or financial stability papers).¹¹ The final data set (hereafter referred to as the IDEAS/RePEc Data Set) went through numerous controls to ensure its completeness and correctness. Nevertheless, it was not possible to extract and match all information. For example, it is not possible to match the author to his or her affiliation if the name listed in the paper differs from the name registered in the database (if it has changed due to marriage, for example). It is also not possible to match a working paper to its journal version if the title has changed significantly. Moreover, not all the harvested information is reported for each research paper. For instance, abstracts are assigned to more

⁹Five central banks in Europe which report some characteristics of their research but do not have any public research publication series, or the series is very short, or the publications are not in English, were omitted from the second data set.

¹⁰The data were extracted from the IDEAS/RePEc database during the first half of 2020. As such, it represents the information which applies to the date on which it was collected. For example, the author's affiliation represents the affiliation reported by her at that time. As such, it does not represent a historical record of the affiliations to which the author reported at the time of paper's publication. A scientific journal is any publication series referred to as a journal by IDEAS/RePEc. As such, it includes central banks' in-house journals (e.g., the Federal Reserve Bank of Richmond's Economic Brief, FRBSF Economic Letter, Bulletin de la Banque de France, or Visnyk of the National Bank of Ukraine) that represent less than 1 percent of all journal publications.

¹¹The publication series were selected based on their research character so as to create a more or less homogeneous group, which means that rather short and analytical publication series were not included. We also excluded in-house journals, as these are by their nature closer to scientific journals than working paper series. Moreover, if the working paper (discussion paper, occasional paper, etc.) is then published in the central bank's in-house journal, it is indicated in our sample.

than 99.6 percent of publications, keywords to about 83 percent, and JEL codes to about 78 percent. About 69 percent of all authors have an assigned affiliation. In addition, the data are subject to the completeness and quality of reporting by central banks, listed journals, and authors themselves. A few central banks in Europe either do not report their research publications in the database at all or do not have such a series. Therefore, some information was harvested manually where available to complete the coverage of European countries (see Table A.2 in Appendix A). We excluded entries with non-English abstracts (only about 2.3 percent) and we expanded the IDEAS/RePEc Data Set with some additional information on impact factor and the gender of the authors. Gender is identified using the R package *gender* based on the authors' first name (85 percent of all first names are recognized and assigned a gender with an average probability (precision)¹² of 98 percent). Impact factors, as reported by the IDEAS/RePEc database, are assigned both to the central banks' research publication series and to the scientific journals in which some of the research papers are published.

3. How to Understand and Organize Research in Central Banks

The views on what constitutes research may vary—both across different central banks and even among the employees of a bank. For some, “research” can be virtually any longer-term project intended to contribute to the work of the bank. For others, the term describes the work that is published, either in a central bank series or academic journal. Of course, in some institutions, research can have also negative connotations as a project with little or no relevance for addressing any practical policy questions (Meyer et al. 2008).¹³ For the purpose of this paper, the term research refers to the systematic study of a question using scientific methods aimed at expanding

¹²The gender is inferred using historical data sets of first names and dates of birth. In this way it is possible to report the probability that a name was male or female. For more details, see the documentation on the package.

¹³Meyer et al. (2008) in their report noted that “the Committee [the External Review Committee of economic research activities at the Bank of Canada] discovered that the term ‘research’ has a very wide range of meanings for many senior managers as well as staff economists at the Bank.”

the general state of knowledge, with more or less direct relevance to policymaking.¹⁴

Both basic and applied research should be present in central banks, as they complement each other. Basic research is generally tasked with thoroughly investigating a problem to push back the frontiers of knowledge. Applied research, on the other hand, is supposed to have an immediate use and visible value added. Although basic research may not have any immediate application, it may have positive long-term spillovers for the institution and for the economy as a whole. One type of research cannot (or should not) live without the other. Basic research provides fundamental knowledge that applied research can build upon. Likewise, basic research benefits from specific improvements and new fundamental questions generated by applied research. In the context of central banking, basic research is motivated by a desire to increase understanding of the fundamental principles grounding the institution. Such research can try to understand and review, for example, the central bank's role in society, its contribution to public welfare, its aim and impact, or the functioning of the whole economy, including financial institutions and financial markets. Applied research, on the other hand, is usually designed to answer a specific research question aimed at an immediate application. This can include, for example, assessing the transmission of a particular monetary, macroprudential, or supervisory policy instrument, or developing and refining a particular methodology, model, or indicator.

Research in central banks usually stands somewhere between basic and applied research. Policymakers appreciate frontier thinking but are generally more interested in research that addresses practical issues faced by the central bank. This is where the two types of research may clash. With basic research, value added is mainly assessed by some form of publication metric, such as the impact factor of the journal publication. However, applied research tends to


¹⁴It remains an open question how to draw a line between short-term analytical tasks and research projects. If we adopt the perspective that research is eventually meant to be published, analytical tasks addressing specific policy questions with no expectation of generating any publishable research output would not be viewed as research. Nevertheless, we highlight that both analytical tasks and research are intertwined and complementary. Analytical tasks are typically informed by current research and sometimes can spur new research topics themselves.

address aspects of less interest to journal editors and may not even be meant for journal publication in the first place. This is apparent from Figure 1, which shows that the average share of central banks' research papers published in scientific journals is close to 40 percent and rising only slightly over time (see Figure 1, panel A).¹⁵ On top of that, the heterogeneity among central banks is quite pronounced (see Figure 1, panel B). Assessing the merits of applied research is thus much more difficult. One cannot simply apply academic criteria to central bank research publications. Many central banks use their own publication series to exchange ideas about bank-specific topics with other central banks and regulatory institutions, academia, or the wider society. They can serve as a strong conceptual and empirical basis for conducting monetary policy and maintaining price and financial stability but research publications also help to increase central banks' transparency and sometimes even autonomy. Mudge and Vauchez (2016) argue that the ECB used its research and academic credentials to affirm its leadership in the Eurosystem and build its capacity to work autonomously from national central banks.

In-house research in central banks has value added along multiple dimensions. First, there are the *direct* benefits. Research can aid policymakers in dealing with current policy issues, such as determining the driving factors of current phenomena, analyzing the impact of measures taken, and quantifying the implications of alternative policy choices. Under the same category, research is well suited to deal with future policy issues as well. Second, there are plenty of *indirect*, often overlooked, benefits. High-quality research publications can enhance a central bank's reputation and increase its credibility and make it easier to defend policy actions both publicly and privately (Issing 2005; Dincer and Eichengreen 2007; Blinder et al. 2008). A modern central bank strives to be predictable to the market most of the time and extensively communicates its decisions to the public. Moreover, having high-quality research is self-sustaining, as the central bank has a greater chance of attracting (and keeping) high-quality economists. These benefits are important to maintain, even at the cost of letting some research time be spent on topics outside the central bank's mandate.

¹⁵The noticeable decline in working papers published in journals since 2014 is due to publishing delays (Björk and Solomon 2013).

Figure 1. Share of Central Banks’ Research Papers Published in Scientific Journals




Note: Panel A: The total number of research papers published by the 50 central banks in their own publication series between 2000 and 2019 is 21,908; the total number of central banks’ research papers published in scientific journals is 7,455. Panel B: Each point of the distribution refers to the share of central banks’ research papers published in scientific journals in the given year between 2000 and 2019. In four instances the share is 100 percent, with the number of working papers produced being only one or two in a given year and for a given central bank.

There is a high degree of heterogeneity as to the amount of research activities carried on by individual central banks. The research activities performed by central banks are not limited to publications only. Figure 2 depicts the degree of engagement in various research activities, summarized in the *overall research score*, for all the national central banks in our sample (see Section 2). The maximum score is 17 (Bank of England) and the minimum is 2 (the central banks of Romania and the Republic of Moldova). The difference is especially pronounced between central banks in Eastern and Western Europe, with a higher *overall research score* achieved in the West (see Figure 2). Nevertheless, there are three exceptions—the Czech National Bank,¹⁶ the Central Bank of Hungary, and the Bank

¹⁶For more details on the Czech National Bank’s research model, see Malovaná (2020).

Figure 2. Overall Research Scores of Central Banks—United States and Europe



Note: The information was collected manually using the central banks’ websites and social media accounts. The detailed data behind the scores can be found in the working paper version of this article (Malovaná, Hodula, and Rakovská 2020) and in Table A.1. The choropleth maps do not reflect the scores for the Federal Reserve Board (FRB), the European Central Bank (ECB, including the ESRB), and the Bank for International Settlements (BIS).

of Lithuania, which all exhibit an above-average *overall research score*. Even though each of these three central banks achieves its high score by different means (i.e., by engaging in a different set of research activities), all of them are catching up quite well with the well-established and more research-oriented central banks in Europe. Unsurprisingly, all the Federal Reserve Banks achieve a high *overall research score* of between 12.5 and 16.75, reflecting the high importance assigned to research in the system. All these institutions are recognized for their intense investment in scientific prestige and scholarly research.

The *overall research score* consists of three sub-indices—a *research organization score*, a *research publication score*, and a *research popularization score*. Each of the following subsections focuses on one area (for more details on the sub-indices, see Appendix A).

3.1 Research Organization in Central Banks

The vast majority of the central banks in our sample have a separate research department.¹⁷ Most of these departments (divisions, units) fall under the area of monetary policy, forecasting, or general economic analysis. Even though the majority of the central banks in our sample (28 out of 55) have a legal mandate to maintain financial stability, only a few of them have reflected this in the organizational structure of their research activities. Only eight central banks have their own financial research department—three among the U.S. Federal Reserve banks; four among the central banks of the E & SEE, BC region; and the ECB.

Some central banks have affiliated research institutes, which usually focus on a special economic topic (for example, the Bank of Finland Institute for Economies in Transition or the Consumer Finance Institute at the Federal Reserve Bank of Philadelphia). The purpose of such affiliated institutes is to conduct high-quality research and to act as a center of excellence for the selected area. The activities of affiliated centers are usually performed by researchers drawn from across the central bank, but there are also some institutes formed of employees from academia only (the Einaudi Institute for Economics and Finance created by the Bank of Italy) and institutes with a mixture of the two (the Center for Excellence in Finance and Economic Research at the Bank of Lithuania).

3.2 Publication Activities of Central Banks

Central banks publish their research outcomes in various publication series. The most prevalent form of research publication series are working paper series (28 out of 40 central banks in Europe and all 12 Federal Reserve Banks), discussion paper series (8 central banks in Europe), and occasional paper series (5 central banks in Europe). In addition, many central banks also use more specific series, such as technical publications (Research Technical Papers—Central Bank of Ireland), policy-related publications (Research and

¹⁷There are six central banks for which we report no research unit, five of them located in the E & SEE, BC region (Bosnia and Herzegovina, Croatia, Republic of Kosovo, Romania, and Ukraine) and one, the Central Bank of Iceland, in the W & NWE region.

Policy Notes—Czech National Bank), or finance-related publications (Financial Stability Papers—Bank of England; Banking and Financial Studies—Deutsche Bundesbank). An overview of central banks' publication series can be found in Table A.2 in Appendix A.

The *research publication score* summarizes information not only on the existence of such a series but also on open-source research data and the presentation of central banks' publication activity via its own website and via IDEAS/RePEc. A number of central banks publicly share their research data,¹⁸ which then creates a positive externality, as it allows other researchers to replicate the results or to follow up on that research. Nevertheless, these activities are dominated by the U.S. Federal Reserve Banks (11 out of 12), while there are only 5 central banks in Europe that share their research data publicly.¹⁹ Next, a number of central banks promote successful placement of their papers in peer-reviewed journals on their websites. This tendency, however, is recognized mainly across the U.S. Federal Reserve Banks.²⁰

3.3 *Research Popularization in Central Banks*

Central banks are dedicating more and more resources to communicating and enhancing the visibility of their research. Popularization is an important part of overall research activities, as it connects the central bank and its researchers to the wider society. There are a number of ways of increasing the visibility of research outcomes; we have identified three sets of such activities within central banks: (i) use of social media, (ii) a separate webpage (or section of a website) dedicated to research, and (iii) additional forms of presentation or communication, such as blogs, newsletters, and bulletins.

First, central banks are becoming more active on social media. Almost all central banks in Europe (except for the Central Bank of Montenegro) and all the U.S. Federal Reserve Banks use at least

¹⁸We consider here only research data that can be shared publicly. Naturally, some data used in central bank publications cannot be publicly shared due to confidentiality reasons.

¹⁹Denmark, Germany, Lithuania, the Netherlands, and the United Kingdom.

²⁰Even though almost all the central banks in our sample successfully placed at least one research paper in a scientific journal, only 30 percent of central banks in Europe report that on their websites.

one social media platform as a modern channel of communication. The most frequently used social media among central banks are Twitter and LinkedIn. The majority of these central banks even have their own YouTube channel. Many of these central banks use social media for research popularization. This trend is more common and more frequent among the U.S. Federal Reserve Banks and larger central banks in Europe with a well-established research tradition. Moreover, five U.S. Federal Reserve Banks (New York, Philadelphia, Richmond, St. Louis, and Cleveland) and two national central banks in Europe (the Bank of England and the Bank of Finland) have established a separate Twitter account exclusively for posts about research work. All in all, the U.S. Federal Reserve Banks use social media for research purposes much more than European central banks. We find the lowest share in the E & SEE, BC region, where only 22 percent of banks advertise their research work on social media, while the share is 92 percent for the U.S. region (including the Federal Reserve Board). Table A.3 in Appendix A further shows a very similar pattern for the intensity with which central banks that use their social media for research purposes utilize their accounts in general (e.g., the number of Twitter tweets or the number of years since joining the media as a whole). The situation changes when one examines the popularity of these central banks' social media profiles. For example, central banks in the W & NWE region, as well as ECB & BIS, have on average a larger number of Twitter and LinkedIn followers compared to U.S. banks, and surprisingly, all the European regions outrun the United States in the average number of YouTube video views.

Second, central banks are leveraging the potential of their websites to increase the visibility of their in-house research. The majority of central banks in Europe (80 percent), all the U.S. Federal Reserve Banks, the ECB, and the BIS have a separate webpage (or section of a website) dedicated to research activities.²¹ They use this website to highlight research topics prioritized by the central bank (more than 50 percent of the central banks in our sample),²² to promote

²¹Seven banks in the E & SEE, BC region and one central bank in the EA orig. region have no research website.

²²More than 67 percent of central banks in Europe engage in such activity, while only 42 percent of the U.S. Federal Reserve Banks do so.

research events such as conferences and workshops (94 percent), to offer some form of fellowship or traineeship (60 percent), and to communicate opportunities for external collaboration (64 percent). Central banks are less active when it comes to the organization of research competitions and awards (20 percent, mostly from the E & SEE, BC region) and calls for projects (14 percent). Central banks also promote these activities via social media or conventional media. Another way to increase the visibility of research is via individual researchers' profiles, which provide a more personalized view. Central banks may use this tool to boost the scientific prestige of their research staff. All the U.S. Federal Reserve Banks and the majority of central banks in Western Europe (including the ECB and the BIS) promote their researchers on their websites. However, less than one-third of central banks in the E & SEE, BC region follow this practice.

Third, some central banks are exploiting the potential of additional forms of presentation and communication, such as research blogs, bulletins, and newsletters. These types of presentation rarely offer new information, but they do summarize existing research outcomes and activities in a more "digestible" (i.e., shorter, simpler, and non-technical) form. The use of these additional forms of presentation and communication of research outcomes is concentrated in the United States and Western Europe. Only about one-third of central banks in the E & SEE, BC region use one or more of them. Research newsletters are the most common among these three (employed by 55 percent of central banks in our sample). However, they take the form usually of a simple e-mail notification about website updates (67 percent) and less frequently of an original document (23 percent are e-mail campaigns; 10 percent are new research publications). Research blogs and research bulletins are less common. Only about one-third of central banks in the United States and Western Europe use research blogs, and about one-quarter of these central banks publish research bulletins.

4. Research in Central Banks: Some Stylized Facts

There is a significant heterogeneity in research publication activity among central banks from different regions. Table 2 presents

Table 2. Summary Statistics—IDEAS/RePEc Data Set

	EA Orig.	W & NWE	E & SEE, BC	USA	ECB & BIS	Total
No. of CB Research Papers	6,692	1,880	1,982	8,161	3,193	21,908
No. of Journal Publications	2,147	520	399	3,252	1,137	7,455
No. of Authors	4,741	1,512	1,577	4,069	1,433	13,332
<i>of which:</i>						
Female Authors	1,201	316	429	867	307	3,120
Male Authors	2,586	936	724	2,653	894	7,793
Registered in IDEAS/RePEc	2,253	799	713	2,206	621	6,592
<i>of which:</i>						
With Multiple Affiliations	620	155	184	514	158	1,631
With Affiliation to CB	1,008	471	371	727	276	2,853
With Affiliation to CB and at Least One Other Affiliation	197	76	112	187	34	606

(continued)

Table 2. (Continued)

[illegible]


basic summary statistics for the IDEAS/RePEc Data Set used in the rest of this section (see Section 2). The U.S. Federal Reserve Banks, including the FRB, achieve the highest figures in most categories, reflecting the important role of research in the U.S. system. The share of research papers published by the U.S. Federal Reserve Banks and the FRB in scientific journals between 2000 and 2019 is more than 40 percent. Central banks in the original EA member states and supranational institutions (the BIS and the ECB, including the ESRB) have shares about 10 percentage points lower. The U.S. region has the highest number of publications per author, reflecting higher productivity and/or an opportunity to devote more time to research. Fairly similar “productivity” can be observed in the BIS and the ECB. The U.S. region shows a relatively high share of authors with multiple affiliations, which translates to stronger collaboration networks (see Subsection 4.2) with positive synergies for the quantity and quality of research outcomes. However, the male-to-female authors ratio is the second worst among the regions, with 3.4 times more male than female authors (see also Subsection 4.3).

4.1 Topics Researched in Central Banks

In terms of JEL codes, the broad research areas selected by central banks are fairly stable over time, reflecting central banks’ core activities. Unsurprisingly, more than half of the research falls into two categories—*E: Macroeconomics and Monetary Economics* and *G: Financial Economics* (see Figure 3, panel A).²³ This reflects the two main objectives—price stability and financial stability—and the two main sets of instruments—monetary policy and prudential policy—usually assigned to central banks. Even though some central banks may not be directly assigned the goal of financial stability, they have to be interested in the pursuit of this goal, because any disruption to financial institutions and financial markets has a direct impact on the monetary policy transmission mechanism (Smets 2014). The third most common area of research according

²³In contrast, Essers, Grigoli, and Pugacheva (2020) report that almost half of all the IMF working papers published in the period 1990–2017 address research topics connected to *E: Macroeconomics and Monetary Economics* or *F: International Economics*.

Figure 3. JEL Code Analysis




Note: The figure depicts the percentage of publications with a particular JEL code assigned. E: Macroeconomics and Monetary Economics; G: Financial Economics; C: Mathematical and Quantitative Methods; F: International Economics; D: Microeconomics; G2: Financial Institutions and Services; E5: Monetary Policy, Central Banking, and the Supply of Money and Credit; E3: Prices, Business Fluctuations, and Cycles; E4: Money and Interest Rates; G1: General Financial Markets.

to JEL codes is *C: Mathematical and Quantitative Methods*. This reflects the strong rigorous basis of the policy decisionmaking process in central banks, and consequently the need for quantitative research inputs into the policy discussion. A more detailed breakdown of JEL categories shows that the main focus of central banks' research is indeed on financial institutions and financial markets, monetary policy, prices and business cycle fluctuations, and money, credit, and interest rates (see Figure 3, panel B).

Following the GFC, financial stability has emerged as the leading research topic. A simple word cloud of keywords and words in abstracts shows that central banks' research agenda *before* the GFC focused mostly on monetary policy issues and macroeconomic topics in general, such as inflation targeting, exchange rate dynamics, convergence, and unemployment (see Figure 4, upper blue part). Following the outbreak of the GFC in 2008, the topics visibly changed in favor of the macro-finance area (see Figure 4, lower red part). The leading research agenda *after* the GFC is related to financial

Figure 4. Word Cloud Before and After the GFC




Note: Only keywords (the most frequent words in abstracts) with a minimum count of 150 (1,500) are included. The upper (blue) part of the clouds reflects the period 2000–07, while the lower (red) part refers to the period 2008–19. We report word clouds based on both keywords and abstracts because about 17 percent of the publications in our sample do not report keywords, while more than 99 percent report abstracts.

stability policy conduct, with numerous papers analyzing the role of credit dynamics and performing financial risk assessments.

Looking at the dynamics of the five most common keywords, we spot some long-term trends reflecting economic developments and policy challenges (see Figure 5). First, the keywords nicely mirror changes in policy mandates. That is obvious from the increasing dominance of the keywords “financial,” “credit,” and “bank.” These keywords reflect the rising interest of central banks in research topics related to the conduct of financial stability policies. There is an apparent structural break that occurred somewhere around the outbreak of the GFC in 2008. This is when the relative share of *financial* issues overtook that of *monetary* issues. Second, we document a rather stable role of *inflation*-related topics in relative terms. The absolute number of the keyword “inflation” has been gradually rising since 2000, reflecting the fact that the majority of the central banks analyzed operate in fully fledged inflation-targeting regimes.

Figure 5. Five Most Common Keywords in Central Banks' Research Papers




Note: Panel A: Based on 18,256 research papers with keywords published by all European, U.S., and supranational central banks as indicated in Section 2 in period 2000–19. Panel B: Relative share is calculated as the ratio between the frequency of each keyword and the total number of research papers with reported keywords in every given year.

Cumulatively, the first five words appear in 46 percent of all central banks' research papers with reported keywords over the period analyzed.

4.2 Authorship Networks

Research in central banks is heading towards a more collaborative environment. We follow Adams (2012) and consider authorship as a proxy for research collaboration. While the number of publications per author per central bank is decreasing steadily over time (see Figure 6, panel A), the number of authors per publication is increasing (see Figure 6, panel B). What is most striking is the drop in the share of publications with a single author from 40 percent in 2000 to less than 20 percent in 2019. The share of publications with exactly three authors, on the other hand, has more than doubled in the last two decades, reaching 30 percent in 2019. It is also becoming increasingly common to have publications with four

Figure 6. Collaboration Between Researchers




Note: Panel A: Calculated as the number of research publications published by the central bank in the given year divided by the number of authors participating in the research publications in the given year. Panel B: Calculated as the number of publications with the given number of authors in the given year divided by the total number of publications in the given year.

and five or more authors. This is consistent with the general tendency in economic research identified by numerous studies over the last couple of years (see, for example, Card and DellaVigna 2013; Hamermesh 2013; Kuld and O’Hagan 2018; Essers, Grigoli, and Pugacheva 2020).²⁴ The rising number of collaborators is generating positive knowledge spillovers not only to the direct collaborative partners (Azoulay, Graff Zivin, and Wang 2010; Borjas and Doran 2015) but also indirectly to other researchers who are connected to

²⁴Using article information from the top five and top three economic journals, respectively, Card and DellaVigna (2013) and Hamermesh (2013) show that the distribution of the number of authors has shifted steadily rightward. Specifically, Card and DellaVigna (2013) report that the number of authors per paper increased from 1.3 in 1970 to 2.3 in 2012. This is confirmed by Kuld and O’Hagan (2018) on a large sample of the top 255 economic journals. They show that the share of multi-authored papers increased from 50 percent in 1996 to over 75 percent in 2014. Similarly, by studying the history of IMF Working Papers, Essers, Grigoli, and Pugacheva (2020) found that both the number of authors publishing within the IMF and the number of publications issued in the IMF have increased over the last few decades, with the former having the faster pace.

Figure 7. Authors Affiliation and the Role of Leading Researchers



Note: Panel A: The absolute number of authors affiliated either with the central bank or with some other institution (based on the first reported affiliation if there are multiple ones). Panel B: The percentage share of individual authors (researchers) in the total number of authors in our sample in relation to the percentage share of publications authored by these researchers in the total number of publications in our sample.

them within a complex network (Hsieh et al. 2018). Among other things, there is compelling evidence that teaming up is good for individual productivity (Ductor et al. 2014; Ductor 2015; Kong et al. 2019). Although the number of authors involved in central bank research is growing steadily, the importance of a relatively small number of top researchers remains high (see Figure 7, panel B).

Next, we study the collaboration networks at the level of individual institutions. We calculate the bilateral connections between authors based on their reported affiliation in IDEAS/RePEc database and the central bank in which the research paper was published. As such, each edge in our network represents a collaborative connection while the nodes are individual institutions. The width of the edges reflects the total number of collaboration relationships in a given period. For the purposes of our analysis, we divide the sample into five-year non-overlapping periods that allow us to avoid exceptional cases in the evolution of the network and observe changes

over time. To make the network more visually tractable, we divide the resulting analysis into three parts. First, we take a closer look at the collaboration networks formed by central banks in the European region. Second, we focus on the Federal Reserve System in the United States.²⁵ Third, we analyze the cooperation networks that are formed by central banks with institutions other than central banks (e.g., universities, research institutes, and international organizations).


Central banks form enormous research networks. Research collaboration between national central banks in Europe has grown considerably over time. There is a distinct difference in the density of the collaboration network between the first and the second decade analyzed (see Figure 8). Over the first five-year period (2000–04), just eight central banks accounted for 80 percent of all publications, of which only three central banks—the Bank of England, Deutsche Bundesbank, and Banco de España—accounted for 50 percent of all publications. During the subsequent five-year period (2005–09), the cooperation between these institutions intensified, while some formerly isolated central banks joined the research network. The central banks that increased their collaboration with the leading research cluster the most were from Central Europe (Poland and the Czech Republic) and the Baltic region (Estonia and Lithuania).

The density of the research collaboration network increased significantly in the years following the GFC (see Figure 8). This may have been due to several factors combined. First, it may have been the outcome of an increase in knowledge integration, especially in the new EU member states that joined the European Union (EU) in 2004, 2007, and 2013. The European System of Central Banks (ESCB), which consists of all EU national central banks and the ECB, encompasses many research networks and working groups that are meant to stimulate interaction between researchers and promote information exchange.²⁶ Second, the GFC, among other

²⁵Supranational institutions (ECB and BIS) and their connections to national central banks in individual regions are depicted separately in Figure B.2 in Appendix B.

²⁶Past and present examples of such platforms include the Household Finance and Consumption Research Network (HFCN), the Euro Area Business Cycle Network (EABCN), the Wage Dynamics Network (WDN), and the Macroeconomic Research Network (MaRs). See “Research Networks” on the ECB’s website.

Figure 8. Collaboration Networks Between Central Banks—Europe



Note: Each edge between two central banks (nodes) represents authors affiliated with one of the two central banks whose paper was published in the other central bank. The width of the edge reflects the total number of such authors. As such, the width of the edges does not reflect the total number of research papers between the two central banks but rather the total number of collaboration relationships. Central banks shown in red account for 80 percent of publications (dark red for 50 percent and light red for the additional 30 percent).


BoE: Bank of England; CNB: Česká Národní Banka; BdI: Banca d'Italia; BdE: Banco de España; BdF: Banque de France; SR: Sveriges Riksbank; NB: Norges Bank; SP: Suomen Pankki; BdP: Banco de Portugal; CBol: Central Bank of Ireland; ON: Oesterreichische Nationalbank; DNB: de Nederlandsche Bank; SN: Schweizerische Nationalbank; DB: Deutsche Bundesbank; LiB: Lietuvos Bankas; LaB: Latvijas Banka; EP: Eesti Pank; NBP: Narodowy Bank Polski; NBS: Národná Banka Slovenska; MNB: Magyar Nemzeti Bank; BoG: Bank of Greece; CBoC: Central Bank of Cyprus; CBCG: Centralna Banka Crne Gore; NBoU: National Bank of Ukraine; NBS: Narodna Banka Srbije; NBnRSM: Narodna Banka na Republika Severna Makedonija; SI: Sedlabanki Islands.

things, revealed a dark side of highly connected financial markets (Stiglitz 2010; Haldane and May 2011; Acemoglu, Ozdaglar, and Tahbaz-Salehi 2015; Acharya et al. 2017). The fact that dense interconnection serves as a mechanism for the propagation of shocks is now widely acknowledged. As a result, researchers in central banks have begun to form more intense scientific networks, which has led to shared knowledge, data points, and even whole databases.

The U.S. Federal Reserve System is a complex collaborative network of individual reserve banks, with a few leading the research activities. The density of the research collaboration network has remained very high throughout the two decades (see Figure 9) with no significant difference before and after the GFC. The U.S. network is much richer and more evenly distributed than the European network. In the United States, we do not identify any “lost sheep,” as all of the individual Federal Reserve Banks and the FRB have established strong collaborations with the rest. In terms of the authorship counts, we can identify a cluster of four U.S. central banks that can be (historically) considered the leading ones, accounting jointly for 50 percent of the total publications—the Federal Reserve Bank of San Francisco, the Federal Reserve Bank of Chicago, the Federal Reserve Bank of New York, and the Federal Reserve Board.


The U.S. Federal Reserve Banks and the national central banks in Europe maintain fruitful collaborations with academia and other research institutions (see Figure 10). Most of this research collaboration is region specific, meaning that central banks usually collaborate with institutions in their own countries or regions (such as national universities). A relatively small set of institutions collaborate with central banks in all four regions (in dark blue) or in three out of the four regions (in light blue). The institutions with the most authored publications with central banks are renowned research institutes and universities (see Figure 11). In the United States, these are the National Bureau of Economic Research (NBER), the Institute of Labor Economics (ILE), the Centre for Economic Policy Research (CEPR), and the University of Pennsylvania (Penn). Another of the most active research collaborators with the U.S. Federal Reserve Banks is the International Monetary Fund (IMF). Other academic collaborators hail mostly from U.S. universities. Surprisingly, there are also some important links to universities outside the United States, namely the Università Commerciale Luigi Bocconi

Figure 9. Collaboration Networks Between Central Banks—United States



Note: Each edge between two central banks (nodes) represents authors affiliated with one of the two central banks whose paper was published in the other central bank. The width of the edge reflects the total number of such authors. As such, the width of the edges does not reflect the total number of research papers between the two central banks but rather the total number of collaboration relationships. Central banks shown in red account for 80 percent of publications (dark red for 50 percent and light red for the additional 30 percent). FRB: Federal Reserve Board (Board of Governors of the Federal Reserve System); FRBA: Federal Reserve Bank of Atlanta; FRBB: Federal Reserve Bank of Boston; FRBCH: Federal Reserve Bank of Chicago; FRBC: Federal Reserve Bank of Cleveland; FRBD: Federal Reserve Bank of Dallas; FRBKC: Federal Reserve Bank of Kansas City; FRBM: Federal Reserve Bank of Minneapolis; FRBNY: Federal Reserve Bank of New York; FRBP: Federal Reserve Bank of Philadelphia; FRBR: Federal Reserve Bank of Richmond; FRBSF: Federal Reserve Bank of San Francisco; FRBSL: Federal Reserve Bank of St. Louis.

Figure 10. Collaboration Networks Between Central Banks and Other Institutions



Note: Each grey node represents one institution other than the central bank with which the authors in our sample are affiliated. Central banks are divided into four regions and represented by dark red nodes. Each edge represents authors affiliated with an institution other than the central bank whose paper was published in the central bank in the given region. The width of the edges does not give any information in this case. Central banks shown in red account for 80 percent of publications (dark red for 50 percent and light red for the additional 30 percent). Institutions with authored publications with central banks in all four regions are shown in dark blue; institutions with authored publications with central banks in three out of the four regions are shown in light blue.

and the National University of Singapore. In Europe, most authorship relations are found for the CEPR, while the rest are top-shelf European universities. The CEPR is also the single most important collaborator for all four regions considered.

As pointed out by Claveau and Dion (2018), the relationship between researchers from central banks and those from other research institutions can be considered a two-way stream. On the one hand, central banks actively invite other researchers to collaborate on policy-related topics, which may be pursued either via various visiting research programs or by direct cooperation with these institutions. On the other hand, researchers in central banks are entitled to spend some time on basic research in addition to applied (policy-oriented) research. This allows them to collaborate on topics which may be of greater interest to academics.

Looking at the authorship network connections from a different angle may give us a bird's-eye view of researchers' mobility. As mentioned above, the networks presented are formed of collaborative connections between central bank in which the research paper was published and the authors based on their affiliation. The affiliations reported in the IDEAS/RePEc Data Set usually reflect the *current* institutional assignment of the authors, as they are directly from the authors' profiles on the IDEAS/RePEc website. However, the publication itself is historically linked to a given central bank. As such, the affiliations recorded in this publication are fixed, regardless of the authors' mobility across institutions. For most research papers, it can be expected that at least one of the authors was affiliated with the given central bank at the time of publication. Therefore, if the reported affiliations of all the authors differ from the publishing central bank, it is very likely that at least one author has changed her affiliation since then. Panel A in Table 3 provides the numerical representation of network graphs presented in this section.²⁷ In contrast, panel B employs the approach which disregards the information about the publishing central bank and takes into account only information about authors' affiliations. Comparing panel A and panel B gives a rough estimate of researchers' mobility between central banks and from central banks to other research institutions.

²⁷If the author reports affiliations to multiple institutions, a relationship pair is created for each of the institution–central bank pair.

Even though it is not possible to track each individual movement of researchers between central banks and to other research institutions, this comparison provides at least a simplified view of this trend.


The distribution of collaborative connections calculated based on the two approaches differs between regions, indicating researchers' mobility to the ECB, the BIS, and academia. Taking into account the publishing central bank (the first approach), the most network connections exist between authors from the group of original euro area member states, followed by the ECB and BIS and connections to other research institutions (see Table 3, panel A). These make up about 70 percent of all network connections.²⁸ By the nature of the calculation, the second approach yields fewer network connections on aggregate than the first one. However, we can see that the absolute number of connections has increased significantly for the ECB & BIS–ECB & BIS bucket, and also for some collaboration pairs with other research institutions (see Table 3, panel B). This most likely reflects higher mobility to these institutions. In particular, the result suggests that researchers tend to leave national central banks for supranational institutions and for research institutions other than central banks.²⁹

Finally, we complement the networks formed based on researchers' self-reported affiliation with a network looking at the collaborations based only on publishing central banks. Figure 12 shows a collaboration network between individual authors—nodes. The size of the node reflects the number of collaborative relationships of each author, i.e., the number of other authors with whom she collaborated on her research papers. Authors (nodes) are divided into five groups based on the region of the central bank which published

²⁸Twenty-seven percent for the EA orig.–EA orig. bucket, 12.1 percent for the ECB & BIS–ECB & BIS bucket, and 29.8 percent for all connections with other research institutions.

²⁹The mobility towards these institutions is also apparent from higher relative shares of network connections (the numbers in parentheses in Table 3). The relative share is significantly higher in panel B than in panel A in the ECB & BIS–ECB & BIS bucket and the ECB & BIS–EA orig. bucket, and for most connections with other research institutions. Besides researchers' mobility, the difference between the two approaches may also reflect researchers affiliated with the ECB & BIS or other research institutions publishing their papers in national central banks without the collaboration of researchers from national central banks. However, we consider this highly unlikely or limited in scope.

Figure 12. Collaboration Networks Between Authors



Note: Each node represents one author in our data sample. The size of the nodes reflects the number of relationships of the author, i.e., the number of other authors with whom she collaborated. Authors (nodes) are divided into five groups based on a central bank that published the paper they collaborated on. Authors shown in dark red published mainly in ECB or BIS; authors shown in light red published mainly in U.S. central banks; authors shown in dark blue published mainly in central banks from EA orig. region; authors shown in middle blue published mainly in central banks from W & NWE region; authors shown in light blue published mainly in central banks from E & SEE, BC region. For ease of exposition, only authors with a certain number of publications are shown in the figure (2000-04: at least 3 publications; 2005-09: at least 5 publications; 2010-14: at least 15 publications; 2015-19: at least 9 publications). Each edge between two authors (nodes) represents publications between the two authors. The width of the edges reflects the number of publications between the two authors.


the paper. The width of the edges connecting individual nodes then reflects the number of publications between each pair of authors.

A few paragraphs earlier, we highlighted a significant contribution of a relatively small number of top researchers to the final population of research papers (Figure 7). Such evidence can point to the presence of a small world effect (Goyal, Van Der Leij, and Moraga-González 2006) or a superstar effect (Hsieh et al. 2018). In the small world, a few key (“superstar”) researchers are shown to bridge the gap between institutions. In particular, Hsieh et al. (2018), when defining a “superstar” researcher, takes into account the spillover effects of one researcher on others in a collaboration network. Figure 12 shows that the network between researchers in our data set becomes denser over time, with a few visible clusters in all regions. The importance of a few researchers (in terms of the number of connections to others) is visible, especially in the United States. Over time, a group of researchers with a significant number of connections has also grown in other regions, notably in central banks in the euro area (dark blue) and the ECB and the BIS (dark red). However, from a simple visual inspection, these clusters appear to be concentrated in a single region, which is consistent with Table 3, indicating that most or a significant proportion of interactions between researchers are within a given region.

4.3 Gender Structure of Authorship Teams

Research also continues to be male dominated, but the trend is changing. In absolute terms, men make up about two-thirds of all the authors of the central bank research papers under study (see Figure 13, panel A). However, the share of female authors has risen significantly over time, reaching 25 percent at the end of 2019 (see Figure 13, panel B). This trend is driven by an increasing share of mixed teams (involving both women and men), reflecting greater collaboration among researchers and larger research teams in general (see Section 4.2). Kwiek and Roszka (2021) show that mixed teams tend to be more successful in publishing in more prestigious journals. While Boschini and Sjögren (2007) find that women are less likely to coauthor than men, we are witnessing that the trend is changing (albeit at a slow pace), at least in central banks’ research. Still, all-male research teams continue to dominate, with a more than

Figure 13. Distribution of Authors by Gender




Note: Gender of authors identified based on their first name using the R package *gender* and the related database; see Section 2.

50 percent share of all teams in 2019. Increasing number of studies show that in terms of research performance, the differences between male and female scientists can be explained by the relative skewness of the distribution (Nielsen 2017; Abramo, Aksnes, and D’Angelo 2021). Auriol, Friebe, and Wilhelm (2020) argue that it is important to monitor the gender gap in research, because the topics favored by female economists might deviate from those favored by men. If women are less represented in research teams, then these topics might become systematically under-invested. Furthermore, Zhang et al. (2020) identify significant gender differences in the benefits arising from engagement in international academic collaboration.

Small central banks engage more female researchers. The relationship between the share of female authors and the total number of authors per central bank per year seems to follow a downward-sloping convex curve (see Figure 14). However, this relationship differs between regions. Smaller central banks in Europe (in terms of total number of authors) engage more female researchers. On the other hand, central banks in the United States and in the EA original member states employ a lower share of female researchers. A similar pattern was observed by Auriol, Friebe, and Wilhelm (2020),

Figure 14. Relationship Between the Share of Female Authors and the Total Number of Authors



Note: Gender of authors identified based on their first name using the R package *gender* and the related database; see Section 2.

who showed that the more prestigious the research institution is, the fewer female economists are present. The authors also found that this systematic under-representation is present already at entry level and that the ranking of the institution causes the gap to be even wider.

Central banks rank in the lower half of research and development (R&D) institutional sectors in terms of the proportion of women in research. Table 4 shows that the private non-profit, government, and higher-education sectors exhibit higher shares of female researchers than central banks, while the lowest percentage is recorded by the business enterprise (BE) sector. In 2015–17, women made up more than 40 percent of the research population in the private non-profit, government, and higher-education sectors. These three sectors together employed more than 99 percent of all researchers in the EU as of 2018 (European Commission 2019), which suggests a gradually closing gap. Nevertheless, women are still vastly under-represented in business enterprise research.³⁰ That central banks lie somewhere

³⁰One of the possible drivers of this result is the constant shortage of female graduates (and also doctoral graduates) in the several narrow fields of science,

Table 4. Share of Female Researchers in Different Sectors (%)

	Central Banks	Business Enterprise Sector	Gov. Sector	Higher Education Sector	Private Non-profit Sector
Europe Total	27.0	19.9	40.1	39.6	44.2
EA Original Members	26.7	19.1	38.9	37.6	45.0
W and NW Europe	26.2	20.5	38.4	43.3	42.0
E and SE Europe, and BC	31.7	23.8	43.8	40.5	41.1
2000–04	21.5	18.2	35.5	32.8	40.1
2005–09	25.9	19.2	39.3	41.2	43.5
2010–14	28.2	20.3	41.4	41.2	44.7
2015–17	29.5	21.0	43.8	43.1	49.1
Note: The data cover the period of 2000–17. Unlike the IDEAS/RePEc Data Set, the Eurostat database contains data on only 16 E & SEE, BC countries: BG, CY, CZ, EE, GR, HR, HU, LT, LV, MK, MT, PL, RO, RS, SK, SL. Figure 13 refers to data on the whole sample (including the United States, the ECB, and the BIS); therefore, it may differ slightly from the numbers in this table.					

between the shares for business enterprises and higher education is not surprising. Based on OECD (2015, p. 90), the R&D activities of corporations from the financial sector and its sub-sectors, and hence also central banks' R&D activities, fall under either of the two R&D categories. A growing share of women can be observed also among doctoral graduates, as the average annual growth rate of female doctoral graduates in the EU was 2.4 percent between 2007 and 2016, while that of male doctoral graduates was only 1.4 percent (European Commission 2019).

This comparison should be taken with caution, because the figures for central banks refer to authors of research papers, while those for other sectors represent the headcount of R&D personnel as reported in the enterprise survey.³¹ Nevertheless, by assuming that researchers usually present their results in the form of a research paper, this can be considered a reasonable approximation.

4.4 *Impact Factor*

The impact factor is a standard measure for assessing the academic quality of a scientific publication. In fact, this is much more common for academic publication series than for central banks' publication series. Although we cannot say that central banks' publications with a higher impact factor have higher policy relevance³² (to our knowledge, there is no such indicator), the impact factor can be an important measure of "value for research." Central banks' publication series with higher impact factors are cited more in scientific publications and by other working papers, which indicates the extent to which other researchers find in these series material inspiring their own work.

technology, engineering, and mathematics (STEM) (Huyer 2015; European Commission 2019). Research in these fields covers a significant portion of business enterprise research, so it is reasonable to assume that this broadly discussed issue also affects the gender disparity in science.

³¹The data are gathered via national questionnaires in paper and/or electronic format, and the statistical unit used is enterprise. R&D personnel include all those engaged directly in the R&D activities of the given enterprise (<https://ec.europa.eu/>).

³²By policy relevance we understand, in a broad sense, that a central bank paper becomes one of the inputs into a policy decisionmaking process.


The impact factor³³ of central banks' research publications varies substantially both between and within the regions. Not surprisingly, the ECB, the BIS, and the U.S. Federal Reserve Banks, including the FRB, receive the highest impact factors for their in-house research publication series (see Figure 15, panel A). The recorded dispersion is relatively large in the case of the U.S. region, ranging from as little as 2 to 16.³⁴ When considering only those papers which have been published in scientific journals, we find the U.S. region on top, followed by the two supranational institutions and the W & NWE region (see Figure 15, panel B). The fact that the U.S. Fed has a higher rate of influence on journal publications is not surprising. Historically, American economists contributed by far the largest share of journal publications and were cited much more often than European economists (Frey and Eichenberger 1993; Rybacki and Serwa 2021). In addition, U.S. research institutions and academic journals in economics are of a higher rank (Kalaitzidakis, Mamuneas, and Stengos 2003, 2011) than those located in Europe and, as such, the U.S. academic market is much bigger. Thus, we might be witnessing the effect of both quality and quantity. Central bank researchers can be schooled at top-tier universities with a long tradition of top economic research and have access to a wide range of top-tier academic outlets to publish their papers.

Interestingly, there is a high correlation between the impact factor assigned to the central banks' research publication series and the average impact factor of the papers published in scientific journals (see Figure 16). The impact factors of the U.S. Federal Reserve Banks and the FRB and the ECB and the BIS reflect the high emphasis placed on research in these institutions, the high share

³³We present a recursive impact factor, as reported in the IDEAS/RePEc database. However, when compared to other types of impact factors (simple impact factor and a discounted impact factor), the conclusions remain the same. The simple impact factor is calculated as the number of all citations of papers in a particular series or journal divided by the number of papers in the series or journal (self-citations are not considered). The recursive impact factor is calculated in the same way, except that each citation carries some weight. The discounted impact factor gives more weight to what is cited now. For more details, see Zimmermann (2012).

³⁴The Federal Reserve Bank of San Francisco and the Federal Reserve Bank of New York rank the highest, while the Federal Reserve Bank of Cleveland and the Federal Reserve Bank of Kansas City rank the lowest.

Figure 15. Recursive Impact Factor—Regional Distribution




Note: Panel A: The recursive impact factor assigned to central banks' publication series (working paper series, discussion series, occasional papers, etc.). Panel B: The distribution of the non-zero recursive impact factor assigned to the scientific journals in which the central banks' research papers were published. It accounts for 7,411 papers between 2000 and 2019 in total (data as of March 2020). Distribution with simple and discounted impact factor is similar (available upon request).

of researchers from academia involved in the research, and the related high importance given to publishing in high-impact journals. Another significant factor contributing to the relatively high average impact factor is the fact that research papers on the U.S. economy (and other large economies or international samples) generally have a higher probability of publication in a high-impact scientific journal. It is usually harder to publish small country-specific research or applied research which is of interest to a particular central bank but little interest to an international scientific journal.

The impact factor is an important measure of the scientific performance of academic journals, even though some studies have pointed to its problematic nature (Seglen 1997; Simons 2008; Vanclay 2012). The impact factor is derived from the number of citations received by the papers in a particular series or journal. As such, it may help researchers and others in searching for bibliographic references to a particular paper and enable mutual communication.


Figure 16. Simple Impact Factor—Individual Central Banks



Note: The interquartile range and mean in blue refer to the non-zero simple impact factor assigned to the scientific journals in which the central banks’ research papers were published. The dark red line refers to the simple impact factor assigned to the central banks’ publication series. Horizontal lines separate individual regions as follows: EA orig.; W & NWE; E & SEE, BC; ECB & BIS; and US. Only central banks with publication series with an impact factor assigned in the IDEAS/RePEc database are reported. Recursive and discounted impact factors are reported in Appendix B.

It is also an appropriate indicator of research interest, i.e., the extent to which what is published in a series or journal is valued by future research. However, the informative value of the impact factor may be limited, since a small number of papers published in a series or journal may have received a large proportion of the citations. The overall impact factor is then boosted by this small share of papers and may say a little about the quality of the other papers published in a particular series or journal. This pattern is also visible in our data, which show that a very small proportion of central banks’ research papers published in scientific journals contribute the most to the average and median impact factor (see Figure 17, panel A). The same pattern persists if we divide the sample into individual regions, with a heightened effect in the United States, the ECB, and the BIS (panel B). These central banks are the ones with the highest average impact factors of their in-house publication series.

Figure 17. Importance of Publications in High Recursive Impact Factor Journals



Note: The distribution with simple and discounted impact factor is equivalent (available upon request).

5. Conclusions

In this paper, we explore research in 55 central banks in Europe and the United States, focusing on the quantity and quality of publication activity, topics analyzed, the cooperation networks between individual central banks and with academia, the gender structure of research teams, and research popularization. For this purpose, we employ two unique data sets containing information on different aspects of central banks’ research and more than 20,000 research papers published in these central banks over the 2000–19 period.

There is significant heterogeneity in research activities across central banks from different regions. This heterogeneity is most apparent between larger, well-established central banks with a long research tradition (usually located in the United States and Western and Northwestern Europe) and smaller central banks (usually located in Eastern and Southeastern Europe). The well-established central banks produce more research papers per author, publish a higher proportion of these papers in scientific journals, and generally achieve a higher average impact factor. On the other hand, they employ fewer female researchers. Furthermore, the vast

majority of central banks in our sample have a separate research department, but only a minority distinguish between financial and economic research in their organizational structure. Central banks are also dedicating increasingly more resources to communicating and enhancing the visibility of their research. They have become more active on social media and exploited the potential of additional forms of presentation and popularization, such as research blogs, bulletins, and newsletters. However, the use of social media for research popularization is more common among the U.S. Federal Reserve Banks and larger central banks in Europe with a well-established research tradition.

Taken as a whole, we have identified seven stylized facts of central banks' research. First, financial stability has emerged as the leading topic after the GFC. Based on a word cloud of keywords and abstracts, we find that monetary policy and macroeconomic topics in general dominated before the GFC, while the macro-finance area has taken over since then.

Second, research in central banks is becoming increasingly collaborative. While the number of publications per author per central bank is decreasing steadily over time, the number of authors per publication is on the rise. The most striking is the drop in share of the publications with a single author, from 40 percent in 2000 to less than 20 percent in 2019. The share of publications with exactly three authors, on the other hand, has more than doubled in the last two decades, reaching 30 percent in 2019. Four- or five-member teams are also becoming more common. Partly as a result of the increasingly collaborative environment, the number of papers published in central banks' series has more than doubled over the two decades.

Third, central banks form enormous research networks. The closer collaboration among researchers reflects more intensive cooperation both inside and outside the network of central banks. In other words, the increasing number of authors per central bank publication is due to both a larger number of researchers affiliated with a central bank and a larger number of those affiliated with an institution other than a central bank. Still, we find that most of the research collaboration is region specific, meaning that a central bank is more likely to collaborate with a research institution from its own country or region. Further, our data suggest that researchers tend to

leave national central banks to relocate to supranational institutions, identifying a potential problem associated with human capital flight.

The second and third stylized facts are consistent with the general tendencies in economic research identified by numerous studies over the last couple of years (see, for example, Card and DellaVigna 2013; Hamermesh 2013; Kuld and O'Hagan 2018; Essers, Grigoli, and Pugacheva 2020). They may to a large extent reflect the increasing scale and complexity of central banks' duties, as well as increasing integration, especially among European countries, and increasing interconnectedness of world economies and financial systems. Surprisingly, the increased collaboration with academia has not been reflected in an increased share of central banks' papers published in academic journals. This share stands at nearly 40 percent over the two decades, abstracting from publishing delay.

Fourth, a relatively small share of authors contributes to a relatively large number of central banks' research papers. Specifically, the top 10 percent of authors contributed to about 50 percent of all central banks' publications. This finding suggests (albeit very indirectly) that even central banks may be prone to the "extinction effect" that has been found to be present in academic research (Azoulay, Graff Zivin, and Wang 2010). Specifically, the departure of a leading researcher may result in a decrease in the central bank's publications activity.

Fifth, the share of female researchers increases over time, but the gender gap persists. The proportion of female authors in our sample rises from about 19 percent in 2000 to 25 percent in 2019. This trend is driven by an increasing share of mixed research teams (involving both women and men), while the share of all-women research teams remains stable over time. Further, small central banks engage more female researchers than larger central banks. A similar pattern was observed by Auriol, Friebe, and Wilhelm (2020), who showed that the more prestigious the research institution is, the fewer female economists are present. As for the relative comparison with other R&D sectors, central banks stand somewhere between the business enterprise sector and the government and higher education sectors.

Sixth, larger, well-established central banks with a long research tradition achieve a higher impact factor of their in-house research publication series as well as their papers published in scientific journals. Although we cannot say that central banks' publications with a

higher impact factor have higher policy relevance, the impact factor can be an important measure of “value for research.” Central banks’ publication series that enjoy a higher impact factor serve as a fundamental source of inspiration to other researchers, and as such may have a significant impact on the direction of future research. Nevertheless, the impact factor should be interpreted cautiously, because the overall impact factor—by design—can be boosted by a small number of papers with a high number of citations.

Seventh, a relatively small proportion of central banks’ research papers published in scientific journals contribute the most to the average and median impact factor. This indicates that the pattern described in the previous paragraph may also apply to the impact factor of central banks’ in-house publication series. This may be true because there is a high positive correlation between the impact factor assigned to central banks’ in-house publication series and the average impact factor of the papers published in a scientific journal.

Appendix A. Data

Table A.1. Items Included in the Research Scores

Item	Question	ORS	ROS	RPuS	RPoS
A	Is there a separate research department?	Y	Y	N	N
B	Is financial research separated in the CB's organizational structure?	Y	Y	N	N
C	Is there an associated research institution?	Y	Y	N	N
D	Is there a separate webpage (or section of a website) for research?	Y	N	N	Y
E	Is there a separate webpage (or section of a website) for financial research? (if 0.5, then financial research can be separated from other research—for example, by tags—but there is no separate webpage for it)	Y	Y	N	N
F	Are the research areas promoted on the research website? (if 0.5, then the research areas are stated only briefly)	Y	N	N	Y
G	Is there a webpage which summarizes researchers' profiles?	Y	N	N	Y
H	Is there a working paper (WP) series published by the CB?	Y	N	Y	N
I	Is there a discussion paper (DP) series published by the CB?	Y	N	Y	N
J	Is there any other publication series except for the WP and DP series?	Y	N	Y	N
K	Does the CB publish its own research bulletin?	Y	N	N	Y
L	Is there a research blog in the CB?	Y	N	N	Y
M	Is there a research newsletter in e-mail form? (if 0.25, then the research newsletter takes the form of e-mail notifications about updates to research websites; if 0.5, then the research newsletter takes the form of an e-mail campaign)	Y	N	N	Y
N	Is there a research newsletter in pdf/document form? (if 0.5, then the research newsletter takes the form of a pdf/document)	Y	N	N	Y
O	Have there been any posts related to the research activities of the CB on LinkedIn in the last six months (as of February/March 2020)?	Y	N	N	Y
P	Have there been any posts related to the research activities of the CB on YouTube in the last six months (as of February/March 2020)?	Y	N	N	Y

(continued)

Table A.1. (Continued)

Item	Question	ORS	ROS	RPuS	RPoS
Q	Have there been any posts related to the research activities of the CB on Twitter in the last six months (as of February/March 2020)?	Y	N	N	Y
R	Is there a separate Twitter account for research?	Y	N	N	Y
S	Does the CB have an RePEc account?	Y	N	Y	N
T	Does the CB advertise any calls for projects on its website?	Y	N	N	Y
U	Does the CB offer any form of research cooperation?	Y	N	N	Y
V	Are there any research events organized by the CB?	Y	N	N	Y
W	Is there any webpage which lists the journal articles written by the CB's staff? (if 0.5, then journal articles written by the CB's staff can be accessed only via the researchers' profile pages)	Y	N	Y	N
X	Does the CB provide data sources for its research publications? (if 0.5, then the CB provides research data in general, but these data sets are not explicitly connected to given research publications)	Y	N	N	Y
Y	Does the CB organize any research competitions or appraisals?	Y	N	N	Y
Z	Does the CB offer any trainee/internship program connected to research?	Y	N	N	Y
Note: The full Research Score Data Set is available in the working paper version of this article (Malovaná, Hodula, and Rakovská 2020). ORS—Overall Research Score; ROS—Research Organization Score; RPuS—Research Publication Score; RPoS—Research Popularization Score.					

Table A.2. Publication Series Covered in the IDEAS/RePEc Data Set

Country	Central Bank	Source	Outlet*	Period	No. of Papers
A. Europe					
1	IT Banca d'Italia	IDEAS/RePEc	WP	2000-19	893
2	RO Banca Nationala a Romaniei	Institution's Website	DP	2006-19	537
3	ES Banco de España	IDEAS/RePEc	OP	2002-19	28
4	PT Banco de Portugal	IDEAS/RePEc	WP	2000-19	782
5	MT Bank Centrali ta' Malta	IDEAS/RePEc	OP	2000-19	96
6	UK Bank of England	IDEAS/RePEc	WP	2000-19	382
			WP	2013-19	33
			WP	2000-19	737
			FSP	2006-18	44
			MPCDP	2000-19	51
7	GR Bank of Greece	IDEAS/RePEc	WP	2003-19	276
8	AL Banka e Shqiperise	Institution's Website	RP	2000-19	82
9	SL Banka Slovenije	Institution's Website	WP	2013-19	20
			DP	2000-19	77
10	Banque Centrale du Luxembourg	IDEAS/RePEc	WP	2001-19	124
11	Banque de France	IDEAS/RePEc	WP	2000-19	671
			OP	2006-17	10
12	BG Bulgarian National Bank	Institution's Website	DP	2001-19	61
13	CY Central Bank of Cyprus	IDEAS/RePEc	WP	2007-19	55
14	IE Central Bank of Ireland	IDEAS/RePEc	RTF	2000-19	214
15	CZ Česká Národní Banka	IDEAS/RePEc	WP	2002-19	281
			RPN	2003-19	36
16	DK Danmarks Nationalbank	Institution's Website	WP	2002-19	141
17	NL de Nederlandsche Bank	IDEAS/RePEc	WP	2004-19	660
18	DE Deutsche Bundesbank	IDEAS/RePEc	DP	2013-19	346
			DP1ES	2000-11	368
			DP2BFS	2003-11	120

(continued)

Table A.2. (Continued)

Country	Central Bank	Source	Outlet*	Period	No. of Papers
<i>A. Europe</i>					
19	EE	Eesti Pank	IDEAS/RePec	2000–19	159
20	HR	Hrvatska Narodna Banka	IDEAS/RePec	2000–19	53
21	LV	Latvijas Banka	IDEAS/RePec	2001–19	75
22	LT	Lietuvos Bankas	DP	2009–19	14
			WP	2008–19	66
			DP	2016–19	16
			OP	2016–19	13
23	HU	Magyar Nemzeti Bank	IDEAS/RePec	2000–19	210
24	MK	Narodna Banka na Republika Severna Makedonija	OP	2000–19	72
			WP	2000–19	60
25	SK	Narodná Banka Slovenska	IDEAS/RePec	2006–19	69
26	RS	Narodna Banka Srbije	IDEAS/RePec	2005–15	31
27	PL	Narodowy Bank Polski	IDEAS/RePec	2002–19	294
28	UA	National Bank of Ukraine	IDEAS/RePec	2016–19	9
29	BE	Nationale Bank van België/Banque national de Belgique	IDEAS/RePec	2000–19	334
30	NO	Norges Bank	IDEAS/RePec	2003–19	337
31	AT	Oesterreichische Nationalbank	IDEAS/RePec	2000–19	187
32	IS	Sedlabanki Íslands	IDEAS/RePec	2000–19	76
33	SW	Schweizerische Nationalbank	IDEAS/RePec	2004–19	215
34	FI	Suomen Pankki	RDP BOFIT	2000–19	597 1,064
35	SE	Sveriges Riksbank	WP	2000–19	467 279

(continued)

Table A.2. (Continued)

Country	Central Bank	Source	Outlet*	Period	No. of Papers	
B. United States						
36	Georgia	Federal Reserve Bank of Atlanta	WP	2000–19	494	494
37	Massachusetts	Federal Reserve Bank of Boston	WP	2000–19	310	358
38	Ohio	Federal Reserve Bank of Cleveland	SRAWP	2007–19	48	499
39	Texas	Federal Reserve Bank of Dallas	WP	2000–19	197	570
40	Illinois	Federal Reserve Bank of Chicago	GIWP	2007–19	373	477
41	Missouri	Federal Reserve Bank of Kansas City	WP	2000–19	278	278
42	Minnesota	Federal Reserve Bank of Minneapolis	WP	2000–19	166	166
43	New York	Federal Reserve Bank of New York	WP	2000–19	806	806
44	Pennsylvania	Federal Reserve Bank of Philadelphia	WP	2000–19	657	657
45	Virginia	Federal Reserve Bank of Richmond	WP	2000–19	262	262
46	California	Federal Reserve Bank of San Francisco	WP	2000–19	571	571
47	Missouri	Federal Reserve Bank of St. Louis	WP	2000–19	876	876
48	Washington, D.C.	Federal Reserve Board	FEDS IFDP	2000–19 2000–19	1,537 610	2,147

(continued)

Table A.2. (Continued)

Country	Central Bank	Source	Outlet*	Period	No. of Papers
<i>C. International Institutions</i>					
49	Bank for International Settlements European Central Bank European Systemic Risk Board	IDEAS/RePEc	WP	2000–19	752
50		IDEAS/RePEc	WP	2000–19	2,335
51		IDEAS/RePEc	WP	2016–19	106
Note: *WP—Working Papers; DP—Discussion Papers; OP—Occasional Papers; FSP—Financial Stability Papers; MPCDP—Monetary Policy Committee Discussion Papers; RP—Research Papers; RTP—Research Technical Papers; RPN—Research and Policy Notes; DP1ES—Discussion Paper Series 1: Economic Studies; DP2BFS—Discussion Paper Series 2: Banking and Financial Studies; RDP—Research Discussion Papers; BOFIT—BOFIT Discussion Papers; SRAWP—Supervisory Research and Analysis Working Papers; GIWP—Globalization Institute Working Papers; FEDS—Finance and Economics Discussion Series; IFDP—International Finance Discussion Papers.					


Table A.3. Central Banks in Social Media

	Social Media	EA Orig.	W & NWE	E & SEE, BC	US	ECB & BIS
CBs with Soc. Media Profile	Twitter	100%	100%	83%	100%	67%
	YouTube	91%	100%	87%	100%	67%
	LinkedIn	100%	100%	78%	100%	100%
	At Least One	100%	100%	96%	100%	100%
CBs Using Soc. Media for Research	Twitter	64%	17%	9%	85%	67%
	YouTube	45%	17%	9%	54%	67%
	LinkedIn	36%	17%	9%	69%	67%
	At Least One	64%	33%	22%	92%	67%
Avg. No. of Followers since Joining	Twitter	15.5k	300.9k	0.8k	122.3k	333.1k
	YouTube	4k	3.6k	6.5k	18.7k	29.8k
	LinkedIn	32k	196.8k	13.8k	26.9k	222.6k
Avg. No. of Years since Joining	Twitter	8	11	4	11	10
	YouTube	8	4	11	10	8
Avg. No. of Tweets since Joining	Twitter	7.1k	9k	0.4k	12.2k	14.9k
Avg. No of Views since Joining	YouTube	862.8k	608.4k	3,091.4k	496.3k	7,242.7k

Note: The table presents a selection of intensive and extensive margin statistics on the social media usage by central banks. Social media usage data were collected from central banks’ websites and social media accounts. To obtain information on the use of social media for research purposes, we analyzed posts of individual central banks over the last six months on the three most used social media platforms—YouTube, Twitter, and LinkedIn. Data from separate Twitter accounts set up solely for research purposes by the Federal Reserve Board, five other central banks in the United States (Federal Reserve Banks of New York, Philadelphia, Richmond, St. Louis, and Cleveland), and two central banks in Europe (the Bank of England and the Bank of Finland) are not included. However, all eight central banks are considered to use social media for research purposes, and these statistics reflect the situation on their general social media network profiles in order to achieve cross-sample comparability. All the statistics except for the average number of followers on LinkedIn employ the data that refer to the period of March 2020. The LinkedIn statistic refers to the period of the article’s revision—February 2022. Only central banks that were identified as using social media for research purposes are considered (the list of such central banks can be found in the working paper version of this article (Malovaná, Hodula, and Rakovská 2020)).


Appendix B. Additional Figures

Figure B.1. Research Score Groups



Note: The detailed data behind the scores can be found in the working paper version of this article (Malovaná, Hodula, and Rakovská 2020) and in Table A.1. The choropleth maps do not reflect the scores for the Federal Reserve Board (FRB), the European Central Bank (ECB, including the ESRB), and the Bank for International Settlements (BIS).

Figure B.2. Authorship Networks Between Central Banks—Regionally



Note: Based on authors’ self-reported affiliation in IDEAS/RePEc. The edge width indicates the relative frequency of authorships between the listed central banks (i.e., the number of publications in the given time period). Central banks shown in red account for 80 percent of publications (dark red for 50 percent and light red for the additional 30 percent).

References

Abramo, G., D. W. Aksnes, and C. A. D’Angelo. 2021. “Gender Differences in Research Performance within and between Countries: Italy vs Norway.” *Journal of Informetrics* 15 (2): 101–44.

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi. 2015. “Systemic Risk and Stability in Financial Networks.” *American Economic Review* 105 (2): 564–608.

- Acharya, V. V., L. H. Pedersen, T. Philippon, and M. Richardson. 2017. "Measuring Systemic Risk." *Review of Financial Studies* 30 (1): 2–47.
- Adams, J. 2012. "The Rise of Research Networks." *Nature* 490 (7420): 335–36.
- Auriol, E., G. Friebel, and S. Wilhelm. 2020. "Women in European Economics." Mimeo.
- Azoulay, P., J. S. Graff Zivin, and J. Wang. 2010. "Superstar Extinction." *Quarterly Journal of Economics* 125 (2): 549–89.
- Björk, B.-C., and D. Solomon. 2013. "The Publishing Delay in Scholarly Peer-reviewed Journals." *Journal of Informetrics* 7 (4): 914–23.
- Blinder, A. S., M. Ehrmann, M. Fratzscher, J. De Haan, and D.-J. Jansen. 2008. "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence." *Journal of Economic Literature* 46 (4): 910–45.
- Bolton, P., S. G. Cecchetti, and L. Reichlin. 2014. "External Evaluation of Research at the Banque de France." <https://www.banque-france.fr/sites/default/files/media/2016/10/25/eer-final-report-29.09.pdf> (accessed March 26, 2020).
- Borjas, G. J., and K. B. Doran. 2015. "Which Peers Matter? The Relative Impacts of Collaborators, Colleagues, and Competitors." *Review of Economics and Statistics* 97 (5): 1104–17.
- Boschini, A., and A. Sjögren. 2007. "Is Team Formation Gender Neutral? Evidence from Coauthorship Patterns." *Journal of Labor Economics* 25 (2): 325–65.
- Card, D., and S. DellaVigna. 2013. "Nine Facts about Top Journals in Economics." *Journal of Economic Literature* 51 (1): 144–61.
- Carney, M. 2014. "One Mission. One Bank. Promoting the Good of the People of the United Kingdom." Speech at Cass Business School, City University, London, March 18.
- Cecchetti, S. G. 2002. "Research in Central Banks." https://szgerzensee.ch/fileadmin/Dateien_Anwender/Dokumente/newsletter/September02.pdf (accessed March 26, 2020).
- Charl  ty, P., D. Romelli, and E. Santacreu-Vasut. 2017. "Appointments to Central Bank Boards: Does Gender Matter?" *Economics Letters* 155 (June): 59–61.

- Claveau, F., and J. Dion. 2018. "Quantifying Central Banks' Scientization: Why and How to Do a Quantified Organizational History of Economics." *Journal of Economic Methodology* 25 (4): 349–66.
- Dincer, N. N., and B. Eichengreen. 2007. "Central Bank Transparency: Where, Why, and with What Effects?" NBER Working Paper No. 13003.
- Diouf, I., and D. Pépin. 2017. "Gender and Central Banking." *Economic Modelling* 61 (February): 193–206.
- Ductor, L. 2015. "Does Co-authorship Lead to Higher Academic Productivity?" *Oxford Bulletin of Economics and Statistics* 77 (3): 385–407.
- Ductor, L., M. Fafchamps, S. Goyal, and M. J. van der Leij. 2014. "Social Networks and Research Output." *Review of Economics and Statistics* 96 (5): 936–48.
- Essers, D., F. Grigoli, and E. Pugacheva. 2020. "Network Effects and Research Collaborations." Working Paper No. 20/144, International Monetary Fund.
- European Commission. 2019. "She Figures 2018." Publications Office Luxembourg.
- Fase, M. M., and W. F. V. Vanthoor. 2000. "The Federal Reserve System Discussed: A Comparative Analysis." SUERF Studies No. 10, Société Universitaire Européenne de Recherches Financières.
- Frey, B. S., and R. Eichenberger. 1993. "American and European Economics and Economists." *Journal of Economic Perspectives* 7 (4): 185–93.
- Gilad, S., M. Maor, and P. B.-N. Bloom. 2015. "Organizational Reputation, the Content of Public Allegations, and Regulatory Communication." *Journal of Public Administration Research and Theory* 25 (2): 451–78.
- Goyal, S., M. J. Van Der Leij, and J. L. Moraga-González. 2006. "Economics: An Emerging Small World." *Journal of Political Economy* 114 (2): 403–12.
- Haldane, A. G., and R. M. May. 2011. "Systemic Risk in Banking Ecosystems." *Nature* 469 (7330): 351–55.
- Hamermesh, D. S. 2013. "Six Decades of Top Economics Publishing: Who and How?" *Journal of Economic Literature* 51 (1): 162–72.
- Hsieh, C.-S., M. D. Konig, X. Liu, and C. Zimmermann. 2018. "Superstar Economists: Coauthorship Networks and Research

- Output.” Discussion Paper No. 13239, Centre for Economic Policy Research.
- Huyer, S. 2015. “Is the Gender Gap Narrowing in Science and Engineering?” In *UNESCO Science Report: Towards 2030*, pp. 85–103. UNESCO Publishing.
- Issing, O. 2005. “Communication, Transparency, Accountability: Monetary Policy in the Twenty-first Century.” *Review* (Federal Reserve Bank of St. Louis) 87 (March/April, Part 1): 65–83.
- Kalaitzidakis, P., T. P. Mamuneas, and T. Stengos. 2003. “Rankings of Academic Journals and Institutions in Economics.” *Journal of the European Economic Association* 1 (6): 1346–66.
- . 2011. “An Updated Ranking of Academic Journals in Economics.” *Canadian Journal of Economics* 44 (4): 1525–38.
- Kong, X., M. Mao, H. Jiang, S. Yu, and L. Wan. 2019. “How Does Collaboration Affect Researchers’ Positions in Co-authorship Networks?” *Journal of Informetrics* 13 (3): 887–900.
- Kuld, L., and J. O’Hagan. 2018. “Rise of Multiauthored Papers in Economics: Demise of the ‘Lone Star’ and Why?” *Scientometrics* 114 (3): 1207–25.
- Kwiek, M., and W. Roszka. 2021. “Gender-based Homophily in Research: A Large-scale Study of Man-Woman Collaboration.” *Journal of Informetrics* 15 (3): Article 101171.
- Lebaron, F., and A. Dogan. 2016. “Do Central Bankers’ Biographies Matter?” *Sociologica* 10 (2): 1–37.
- Malovaná, S. 2020. “How to Organize Research in Central Banks: The Czech National Bank’s Experience.” Research and Policy Note No. 1/2020, Czech National Bank.
- Malovaná, S., M. Hodula, and Z. Rakovská. 2020. “Researching the Research: A Central Banking Edition.” Research and Policy Note No. 3/2020, Czech National Bank.
- Maor, M., S. Gilad, and P. B.-N. Bloom. 2013. “Organizational Reputation, Regulatory Talk, and Strategic Silence.” *Journal of Public Administration Research and Theory* 23 (3): 581–608.
- Meyer, L., M. Eichenbaum, D. Gale, A. Levin, and J. McAndrews. 2008. “External Review of Economic Research Activities at the Bank of Canada.” https://www.bankofcanada.ca/wp-content/uploads/2011/05/ext_review.pdf (accessed March 26, 2020).

- Moschella, M., and L. Pinto. 2019. "Central Banks' Communication as Reputation Management: How the Fed Talks under Uncertainty." *Public Administration* 97 (3): 513–29.
- Mudge, S. L., and A. Vauchez. 2016. "Fielding Supranationalism: The European Central Bank as a Field Effect." *Sociological Review* 64 (2): 146–69.
- Nielsen, M. W. 2017. "Gender and Citation Impact in Management Research." *Journal of Informetrics* 11 (4): 1213–28.
- OECD, F. M. 2015. "Guidelines for Collecting and Reporting Data on Research and Experimental Development 2015." Organisation for Economic Co-operation and Development.
- Rybacki, J., and D. Serwa. 2021. "What Makes a Successful Scientist in a Central Bank? Evidence from the Repec Database." *Central European Journal of Economic Modelling and Econometrics* 13 (3): 331–57.
- Seglen, P. O. 1997. "Why the Impact Factor of Journals Should Not Be Used for Evaluating Research." *BMJ (British Medical Journal)* 314 (7079): 498–502.
- Simons, K. 2008. "The Misused Impact Factor." *Science* 322 (5899): 165.
- Smets, F. 2014. "Financial Stability and Monetary Policy: How Closely Interlinked?" *International Journal of Central Banking* 10 (2): 263–300.
- Stiglitz, J. E. 2010. "Risk and Global Economic Architecture: Why Full Financial Integration May Be Undesirable." *American Economic Review* 100 (2): 388–92.
- Taylor, J. B. 1998. "Applying Academic Research on Monetary Policy Rules: An Exercise in Translational Economics." *The Manchester School* 66 (S): 1–16.
- The Economist*. 2014. "How Jackson Hole Became Such an Important Economic Talking Shop." <https://www.economist.com/the-economist-explains/2014/08/20/how-jackson-hole-became-such-an-important-economic-talking-shop> (accessed August 14, 2020).
- Trichet, J.-C. 2007. "The Role of Research in Central Banks and at the ECB." Speech at the award of the Germán Bernácer prize ceremony, Madrid, May 21. <http://www.ecb.europa.eu/press/key/date/2007/html/sp070522.en.html> (accessed March 26, 2020).

- Vanclay, J. K. 2012. "Impact Factor: Outdated Artefact or Stepping-Stone to Journal Certification?" *Scientometrics* 92 (2): 211–38.
- Zhang, M., G. Zhang, Y. Liu, X. Zhai, and X. Han. 2020. "Scientists' Genders and International Academic Collaboration: An Empirical Study of Chinese Universities and Research Institutes." *Journal of Informetrics* 14 (4): Article 101068.
- Zimmermann, C. 2012. "Academic Rankings with RePEc." Working Paper No. 2012-023A, Federal Reserve Bank of St. Louis.

Has Higher Household Indebtedness Weakened Monetary Policy Transmission?*

Gaston Gelos,^{a,b} Federico Grinberg,^a Shujaat Khan,^a
Tommaso Mancini-Griffoli,^a Machiko Narita,^a
and Umang Rawat^a

^aInternational Monetary Fund

^bCEPR

Has monetary policy been less effective since the global financial crisis because of deteriorating household balance sheets? This paper examines the question using household data from the United States. It compares the responsiveness of household consumption to monetary policy shocks in the pre- and post-crisis periods, relating changes in monetary transmission to changes in household indebtedness and liquidity. The results show that the responsiveness of household consumption has diminished since the crisis. However, household balance sheets are not the culprit. More indebted and less liquid households are the most responsive to monetary policy, and their share in the population grew. The decline in the consumption response does not seem to be attributable to households' decreasing interest rate exposure, either.

JEL Codes: E43, E52, E21.

1. Introduction

A common perception among many academics and policymakers is that monetary policy in advanced economies has been less effective since the global financial crisis (GFC) because of higher household debt and associated credit constraints. Amir Sufi summarized this view in 2015 (Sufi 2015): “Monetary policy over the past seven years

*The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. Author contact: International Monetary Fund, 700 19th St, NW, Washington, DC 20431. E-mail addresses: ggelos@imf.org, fgrinberg@imf.org, skhan3@imf.org, tmancinigriffoli@imf.org, mnarita@imf.org, urawat@imf.org.

has been ineffective because it has channeled interest savings and additional credit to exactly the households that are least likely to change their spending in response. The households that would normally spend most aggressively out of monetary policy shocks are heavily indebted or have seen their credit scores plummet, rendering them either unwilling or unable to boost spending.”

To date, however, the issue has—to our knowledge—not been systematically assessed. While a few studies have examined the role of household balance sheets in monetary transmission, they have focused on the pre-crisis period, and have not directly analyzed whether post-crisis debt levels have impeded transmission.¹ These studies suggest that more indebted and less liquid households react more to monetary policy. The argument is that these households run into collateral and liquidity constraints, which monetary policy directly affects (Aladangady 2014, Di Maggio et al. 2017, Cloyne, Ferreira, and Surico 2018, and Flodén et al. 2021 emphasize households’ cash flows; Luo 2017 focuses on households’ default risk). Using aggregate data, Hofmann and Peersman (2017) also find a stronger impact of monetary policy in economies with high private debt. One open question, however, is whether at very high debt levels effects are different. In these cases, monetary easing may do little to alleviate credit constraints, and stimulate consumption (Sufi 2015; Alpanda and Zubairy 2019; Beraja et al. 2019). The responsiveness of households to monetary policy may thus display an inverted U-shaped pattern, rising as debt levels grow below a certain threshold, and declining thereafter.

In this paper we compare the transmission of monetary policy through household consumption in the pre- and post-crisis periods, and ask whether changes therein can be explained by the evolution of

¹Without discussing monetary policy effects, Mian, Rao, and Sufi (2013) and Kaplan, Violante, and Weidner (2014) find that leverage and liquidity significantly affect household’s propensity to consume. Similarly, some empirical studies show adverse effects of high debt on consumption, although they do not examine monetary policy effects (such as Dynan 2012; Drehmann, Juselius, and Korinek 2017; International Monetary Fund 2017; Melzer 2017; Mian, Sufi, and Verner 2017; Alter, Feng, and Valckx 2018). Many studies highlight the adverse effect on aggregate demand from debt deleveraging caused by the housing crisis during the U.S. Great Recession (such as Guerrieri and Lorenzoni 2011; Eggertsson and Krugman 2012; Mian and Sufi 2014; and Eggertsson, Mehrotra, and Robbins 2017).

household balance sheets. To this end, we use quarterly household-level data from the U.S. Consumer Expenditure Survey (CEX) from 1996 to 2014. We first assess average changes in the responsiveness of household consumption to monetary policy shocks, which we identify using exogenous instruments drawn from high-frequency data, in the tradition of Bernanke and Kuttner (2005). We employ both synthetic cohort analysis (which enable us to obtain longer times series and derive local projections) and standard panel data methods that exploit the full micro data set. Next, we explore the role of two household balance sheet variables in driving cross-sectional differences in the responses to monetary policy shocks: indebtedness (mortgage balance relative to house value) and liquidity (liquid assets to monthly income).²

We show that the response of household consumption to monetary policy shocks has diminished since the global financial crisis. We also find that higher-indebted households tend to respond more to monetary policy shocks—particularly with regard to durable consumption—in the pre- and post-crisis periods. While effects appear non-linear, they are not U-shaped, since households with the highest indebtedness respond most to monetary policy shocks. This suggests that household debt did not contribute to lessening the effects of monetary policy over time, since the distribution of debt did not change markedly with the crisis, while its average even increased somewhat.³ We also explore whether the observed reduction in households' interest rate exposure, due to a decline in the share of variable-rate mortgages, has contributed to the weakening of monetary policy transmission to consumption, but fail to find significant evidence to support this notion.

Similar results hold for household liquidity. Households with lower levels of liquid assets react more strongly to monetary policy

²Recent papers suggest that consumption responses to monetary policy should depend on the distribution of households' liquidity; see heterogeneous agent New Keynesian (HANK) models (such as Hedlund et al. 2017, Kaplan, Moll, and Violante 2018, Kaplan and Violante 2018, and de Francisco 2019).

³Justiniano, Primiceri, and Tambalotti (2015) and Yellen (2016) also suggest that debt overhang alone cannot explain the slow recovery from the U.S. Great Recession. Also, Bernanke (2018) does not find strong predictive powers of household balance sheets for economic conditions, although he argues that it does not dismiss the important role of household balance sheets considering the empirical challenges in identifying macro effects.

shocks, both pre- and post-crisis. Again, because the distribution of liquidity across households remained stable over time, liquidity constraints cannot explain the decline in monetary policy effectiveness. The explanation for the lower effectiveness of monetary policy must therefore lie elsewhere, such as in the higher degree of economic uncertainty brought about by the crisis.

2. Hypotheses and Data

The main questions we explore in this paper are as follows:

- (i) Has the response of household consumption to monetary policy shocks declined since the global financial crisis?
- (ii) Do households with greater indebtedness respond more strongly to monetary policy shocks? Is there evidence of non-linearities—in particular, does the responsiveness decline after a certain threshold?
- (iii) Are households with non-fixed-rate mortgages more responsive to monetary policy shocks?
- (iv) Do households with low levels of liquid assets react more to monetary policy shocks? And, are non-linear effects discernable?
- (v) Can shifts in the distribution of household indebtedness and liquidity, and changes in the share of non-fixed-rate mortgages between the pre- and post-crisis periods explain the observed changes in the average response of household consumption to monetary policy?

2.1 *Data: Variables of Interest, Sources, and Summary Statistics*

We use the Consumer Expenditure Survey (CEX)⁴ for household-level consumption, income, and balance sheet data between 1996:Q1 and 2014:Q4. The CEX data are well suited for our analysis for three

⁴CEX data available at <https://www.bls.gov/cex/pumd.data.htm#stata>.

reasons. First, the survey offers rich cross-sectional variation, with about 7,500 households interviewed per quarter. Second, the quarterly frequency is helpful to study the short-run effects of monetary policy on households' consumption. Third, CEX data span a sufficiently long period to compare household behavior before and after the crisis.

We construct measures of durable and non-durable consumption expenditures. This is to allow for the impact of monetary policy to differ across each category of goods, since theory and empirics suggest that the marginal propensity to consume for durable and non-durable goods are significantly different (Souleles 1999; Parker et al. 2013; see Appendix A for more details).

We consider two key characteristics of household balance sheets: indebtedness and liquidity. Indebtedness is defined as the ratio of each household's total mortgage balance (summed over all the properties owned by the household) to the value of the houses it owns, as reported by households. We exclude other liabilities like credit card balances, since fewer households report these and because mortgage debt is the most significant liability for most households.⁵ We define liquidity as the ratio of liquid assets to monthly income, as reported by households. Liquid assets include the total balance on households' checking and savings accounts, and income is after tax. Details are provided in Appendix A.

Table 1 highlights key features of non-durable and durable consumption. On average, households spend four times more on non-durable consumption relative to durable consumption in any given quarter. However, the standard deviation of durable consumption is notably larger than that of non-durable consumption, pointing to the lumpy nature of durable goods purchases (Caballero 1993). Consumption levels differ across housing tenure, especially for durable consumption (see Appendix A), in line with findings in the literature that housing tenure is correlated with consumption decisions (see, for example, Aladangady 2014 and Cloyne, Ferreira, and Surico

⁵The CEX collects mortgage information in all interviews, while it collects other financial information (such as credit card debt) only in the second and fifth interviews. Therefore, we focus on mortgage debt, the largest component of household debt, in examining the effects of indebtedness. Moreover, there does not appear to be any shift in the distribution of credit card debt-to-income ratio post-crisis.

2018). The distribution of consumption quarter-on-quarter growth changes little after the crisis for both durable and non-durable categories, while the distribution of consumption levels shifts slightly to the left after the crisis.

Homeowners take on varying, though generally high, levels of debt (Table 2), in line with findings in Kaplan, Violante, and Weidner (2014) and Hedlund et al. (2017). On average, 74 percent of households own a house, of which almost two-thirds have mortgage debt (see Appendix A).⁶ Average indebtedness among households over the entire sample is high, at nearly 60 percent, as is the standard deviation of consumption, at 42 percent. The distribution is skewed to the right, however, and does not change particularly from pre- to post-crisis, as discussed further in Section 4.

Liquidity levels also vary significantly across households (Table 3). Median liquidity is lowest for renters and highest for homeowners without mortgages. The same is true of standard deviations. The distribution is especially skewed towards lower liquidity levels due to “hand-to-mouth” households whose liquid assets are less than a half of their monthly income (Kaplan, Violante, and Weidner 2014). The share of such households is nearly 60 percent, of which about two-thirds are homeowners which can thus be considered as “wealthy hand-to-mouth” households.⁷ The distribution of liquidity does not change noticeably from pre- to post-crisis, either, as reviewed in Section 4.

2.2 *Identifying Monetary Policy Shocks*

As typical in this literature, we face a tradeoff between overcoming endogeneity and measuring a meaningful relationship between monetary policy and consumption. The former pushes us to seek exogenous monetary policy shocks. However, as these tend to be

⁶More than 80 percent of mortgage contracts in our sample are fixed-rate mortgages.

⁷The share of hand-to-mouth households is likely overstated due to our narrow definition of liquid assets. Kaplan, Violante, and Weidner (2014) find the share to be 31 percent based on a broader definition of liquid assets allowed by granular balance sheet data from the Survey of Consumer Finances for 1989–2010. However, the paper also finds the share of “wealthy hand-to-mouth” households to be around two-thirds, as in our sample.

Table 4. Monetary Policy Shock

	Percentiles			Mean	Std. Dev.	Obs.
	25th	50th	75th			
Full Sample	−0.066	−0.008	0.040	−0.012	0.084	354,685
Pre-GFC	−0.078	−0.007	0.046	−0.011	0.085	233,620
Post-GFC	−0.029	−0.008	0.014	−0.009	0.060	106,277

small, finding a stable and substantial effect on consumption can be difficult.

We identify monetary policy shocks using high-frequency data at the time of monetary policy announcements. We do so in the tradition of Bernanke and Kuttner (2005), by capturing changes in asset prices closely correlated with monetary policy expectations. However, unlike Bernanke and Kuttner, we do not use futures on federal fund rates, since these remained little changed (and close to zero) during the post-crisis period, despite repeated steps taken to loosen monetary policy, such as through quantitative easing (QE) programs.

To find a measure that is equally suitable for pre- and post-crisis periods, we resort to changes in two-year bond yields, taking the cue from Gürkaynak, Sack, and Swanson (2005, 2007) and, subsequently, Gilchrist, López-Salido, and Zakrajšek (2015), Hanson and Stein (2015), and Ferrari, Kearns, and Schrimpf (2017), among others. The identifying assumption is that two-year bond yields on the day prior to a scheduled monetary policy announcement capture market expectations of future policy interest rates, as well as perceptions of policy uncertainty as reflected in term premia. Thus, changes in two-year yields on announcement days reflect the surprise component of monetary policy along both dimensions. We sum monetary policy surprises from all announcements in a given quarter, as in Romer and Romer (2004), to construct measures consistent with our quarterly data on consumption. Table 4 shows that even though the distribution of the identified monetary policy shock has shrunk in the post-GFC period, substantial variation remains, allowing us to use these shocks for assessing monetary policy effectiveness in both pre- and post-GFC periods.

3. Has the Response of Household Consumption Changed Post-Crisis?

Households are only interviewed by the CEX survey for four consecutive quarters, and subsequently drop out of the data set. This limits the assessment of consumption reaction to monetary shocks to a time horizon of three quarters. Therefore, for a first analysis of impulse responses to monetary shocks, we construct synthetic cohorts to obtain longer time series.

Constructing synthetic cohorts amounts to categorizing households at any given quarter according to pre-defined buckets, then linking the data between buckets to create longer time series. The underlying assumption is that households with similar characteristics—belonging to the same bucket—respond similarly to monetary policy shocks. Obviously, such an approach has its limitations, since households can differ along many characteristics which are not controlled for.

We build cohorts using the head of household's birth year and housing tenure. For the grouping by birth year we define 14 groups using five-year birth-year intervals, while for the grouping by housing tenure we retain 3 groups: owners with mortgage, owners without mortgage, and renters. As a result, we build 42 representative consumer units with data for the whole sample period. More details on the construction of synthetic cohort panel data are provided in Appendix B.

We then use the panel of synthetic cohorts to estimate the response of durable- and non-durable consumption to monetary policy, estimating the impulse response function using Jordà's (2005) local projection method:

$$\begin{aligned} \ln \left(\frac{C_{j,t+h}}{C_{j,t-1}} \right) &= \beta_0^{(h)} + \beta_1^{(h)} 2yr_t + \beta_2^{(h)} postGFC \\ &\quad + \beta_3^{(h)} postGFC * 2yr_t + \beta_4^{(h)} X_{j,t} + \beta_5^{(h)} S_t \\ &\quad + \varepsilon_{j,t+h} \quad h = 1, \dots, 12, \end{aligned} \quad (1)$$

where $\ln \left(\frac{C_{j,t+h}}{C_{j,t}} \right)$ is the cumulative log change in real consumption by the synthetic cohort j between periods t and $t+h$, $2yr_t$ is the two-year yield, $X_{j,t}$ is a cohort-specific vector of controls that includes

age and age squared, and S_t is a set of macro controls that includes inflation, GDP growth rate, and quarterly dummies.

To test the hypothesis that the effect of monetary policy has changed after the GFC, we include a dummy variable, labeled *post-GFC* in the above equation, for the post-crisis period (2009:Q1 and onwards) and interact it with the policy rate. The coefficient $\beta_1^{(h)}$ captures the pre-GFC effect of monetary policy and $\beta_3^{(h)}$ captures the additional effect of monetary policy added in the post-GFC. These consumption responses to a contractionary monetary policy are expected to be persistently negative.⁸

We instrument the two-year yield ($2yr_t$) to address endogeneity—the possibility that bond yields reflect monetary policy responses to changes in consumption. As instruments, we adopt exogenous monetary policy shocks from high-frequency data, as discussed earlier. We exploit over-identification to overcome weak instrument bias by using the contemporaneous monetary policy shock and its lags as the instruments. We use the generalized method of moments (GMM) to obtain more precise estimates (see Ramey 2016 and Stock and Watson 2018).⁹


Turning to the results, the pre-GFC effect of monetary policy measured by $\beta_1^{(h)}$ is negative on both durable and non-durable consumption growth, while the additional effect due to the post-GFC $\beta_3^{(h)}$ is positive at most projection horizons (Figure 1).

These results suggest that the responsiveness of consumption to monetary policy has changed and has likely weakened since the crisis. However, the change seems difficult to measure in a precise and robust manner. The size and significance of $\beta_3^{(h)}$ varies as the

⁸Previous studies show that a contractionary monetary policy would generate a hump-shaped drop in consumption and investment in the data (e.g., Christiano and Eichenbaum 2005; Wong 2015; and Cloyne, Ferreira, and Surico 2018), which could be explained by various frictions (e.g., see Christiano, Trabandt, and Walentin 2010; Iacoviello and Neri 2010, and Alpanda and Zubairy 2019). These consumption responses to monetary policy are related to but different from the intertemporal elasticity of substitution (e.g., see Kaplan, Moll, and Violante 2018). For a survey of the estimation of the intertemporal elasticity of substitution, see Thimme (2017).

⁹Results are robust to instrumenting the policy rate with the signal shock and the risk shock described above.

Figure 1. Response of Durable and Non-Durable Consumption to Monetary Policy



Note: GMM estimation, 1996:Q1–2014:Q4. Dependent variable is the accumulated quarterly growth rate in real consumption. Individual data from CEX are aggregated in 42 synthetic cohorts according to housing status and five-year birth-year intervals. In the first-stage regression, the two-year yield is instrumented by monetary policy shocks. All regressions include a constant, aggregate macroeconomic controls (inflation and real GDP growth), and quarterly seasonal effects. Standard errors are robust to heteroskedasticity and autocorrelation. Full line shows the estimated effect, while the dotted lines show the 90 percent confidence interval.

specification of Equation (1) is modified by, for example, changing the set of controls to include more lags. The message we therefore take from this exercise is that it offers suggestive, but not conclusive and precise, evidence for a weakening of monetary policy effects on household consumption in the post-crisis period.

We thus tackle the same question using the full micro data set, without aggregating households in cohorts—at the expense of only

observing consumption growth for three consecutive quarters for any single household. We run the following regression:

$$\ln \left(\frac{C_{i,t+1}}{C_{i,t-1}} \right) = \beta_o + \beta_1 2yr_t + \beta_2 (postGFC \cdot 2yr_t) + \beta_3 postGFC + B\mathbf{Z}_{i,t} + \lambda_{s(t)} + u_{i,t}, \quad (2)$$

where $\ln \left(\frac{C_{i,t+1}}{C_{i,t-1}} \right)$ is the cumulative log change in real consumption for individual household i (instead of a synthetic cohort as above). We focus on two-quarter growth rates in consumption to allow for lags in monetary policy transmission.¹⁰ *postGFC* again denotes a dummy variable that takes the value of 1 since 2009:Q1; $\mathbf{Z}_{i,t}$ denotes household level controls, which include race, education level, age, family size, and marital status. $\lambda_{s(t)}$ stands for seasonal fixed effects.¹¹

We follow the same procedure as before in instrumenting two-year yields ($2yr_t$) using high-frequency monetary policy shocks and their lags as instruments to overcome weak instrument bias, and using GMM estimation.

The results confirm the earlier findings of a weaker impact of monetary policy after the crisis. Overall, we find the expected response of both durable and non-durable consumption to monetary policy shocks. In the pre-crisis period, an expansionary monetary policy shock (a 10 basis point reduction in the two-year yield) increases non-durable and durable consumption by about 3 percent and 2 percent, respectively (Table 5, columns 3 and 4). In the post-crisis period, the response of durable and non-durable consumption to monetary policy is clearly weaker (as seen by positive and significant values of β_2). For durable consumption, the effect is only

¹⁰The rotating nature of data does not allow for a more dynamic analysis of consumption response to monetary policy, an issue we explore using synthetic cohorts. We use two-quarter-ahead consumption growth in the panel analysis to strike a balance between allowing for a transmission lag and not losing too many observations. The results are broadly robust to the choice of one, two, or three quarters.

¹¹See Table A.1 for correlations among consumption growth, household characteristics, and balance sheet variables. Households' balance sheet variables (liquidity and leverage) are not found to be highly correlated with household-level characteristics (family size, education, ethnicity, marital status, etc.).

Table 5. Impact of Monetary Policy on Consumption: Has It Changed After the Crisis?

Variables	Non-durables (1)	Durables (2)	Non-durables (3)	Durables (4)	Non-durables (5)	Durables (6)
Two-Year Yield	-24.24*** (4.46)	-4.37* (2.27)	-26.60*** (4.58)	-17.81** (8.57)	-28.54*** (5.16)	-16.55* (8.76)
Family Size	-0.13 (0.08)	0.06 (0.74)	-0.12 (0.09)	0.07 (0.74)	-0.17* (0.09)	0.04 (0.75)
College Education	0.85*** (0.24)	0.41 (1.97)	0.90*** (0.24)	1.33 (1.88)	0.97*** (0.24)	1.25 (1.91)
White	0.74** (0.29)	2.49 (2.87)	0.71** (0.29)	2.15 (2.84)	0.70** (0.30)	1.86 (2.87)
Married	0.63*** (0.25)	0.90 (2.20)	0.56** (0.25)	0.55 (2.17)	0.51** (0.25)	0.53 (2.20)
Reference Age	0.02*** (0.01)	0.05 (0.07)	0.02*** (0.01)	0.08 (0.07)	0.02*** (0.01)	0.07 (0.07)
Post-GFC			-10.49*** (2.69)	-0.22 (5.56)	-11.31*** (2.85)	0.33 (5.74)
Post-GFC*2_yr Yield			8.54** (3.45)	11.60* (6.50)	9.14*** (3.52)	11.46* (6.79)
Income Growth					0.04*** (0.00)	0.08** (0.03)
Observations	166,921	69,781	166,921	69,781	161,449	68,008
No. of Households	85,246	47,356	85,246	47,356	82,594	46,150
Hansen	Exactly Identified	0.319	Exactly Identified	0.355	Exactly Identified	0.280

Note: GMM estimation, 1996:Q1–2014:Q4. Dependent variable is two-quarter-ahead consumption growth. In the first-stage regression, two-year yield is instrumented by monetary policy shocks. Post-GFC denotes a dummy variable that takes the value of 1 since 2009:Q1. All regressions include a constant, and quarter (seasonal) effects. Clustered standard errors (by households) are reported in parentheses. ***, **, and * indicate statistical significance at the 1, 5, and 10 percent levels, respectively.

marginally statistically significant (Table 5, columns 4 and 6). Our main results are robust to adding income change as an additional control variable (columns 5 and 6).


Household-level controls have a significant and expected impact on households' non-durable consumption. College-educated, white, married, and older households display higher consumption growth following looser monetary policy. However, these characteristics are not found to be important determinants of durable consumption.

When we estimate Equation (2) over the full sample, by removing the post-GFC dummy and its interaction with monetary policy, results show that expansionary monetary policy boosts both durable and non-durable consumption, as expected (Figure 2 and Table 5, columns 1 and 2). The estimated effect is stronger for non-durables; a 10 basis point increase in the two-year yield reduces non-durable consumption by 2.5 percent and durable consumption by 0.5 percent.¹² Results for durable consumption are in general less robust, partly reflecting the diminished response of durables consumption to monetary policy shocks post-crisis, as shown above.¹³ Overall, the effects are comparable to those reported in Wong (2015), who uses the CEX data and finds that a one-standard-deviation expansionary shock (~ 10 – 15 basis points) increases total consumption (both durable and non-durable) by 1.7 percent. The results are qualitatively robust to using monetary policy news shocks based on Nakamura and Steinsson (2018) (see Appendix D for more details).

¹²These results are not directly comparable to those in the literature because this full sample includes the post-GFC period. Moreover, to our knowledge, our novel approach, which uses monetary policy shocks as instruments, has not been used in other studies with micro-level data; the literature tends to use monetary policy shocks as regressors (e.g., Wong 2015) or use other variables to instrument a change in a relevant interest rate (e.g., Aladangady 2017). Further, while Wong (2015) estimates consumption elasticity to monetary policy shocks, we focus on the consumption response to exogenous changes in policy-relevant interest rates, allowing for making more meaningful and policy-relevant conclusions.

¹³Another reason may be the lumpy nature of durables consumption, which implies lags in responses to monetary policy shocks. To partly account for potential lags, we use current and lagged monetary policy shocks as instruments for the two-year yield when estimating Equation (2) for durable consumption growth. The Hansen statistic (test for over-identifying restrictions) shows that instruments are valid. In the case of non-durables, the equation is exactly identified, since we only use the current monetary policy shock as an instrument for two-year yield.

Figure 2. Consumption Response to 10 Basis Point Increase in Two-Year Yield (in percentage points)



4. Does Household Indebtedness Matter?

In this section we ask whether household indebtedness affects the response of consumption to monetary policy impulses. Next, we explore the role of non-linearities, and ask whether the change in the distribution of household indebtedness post-crisis might help explain the lower monetary policy impact on consumption detected earlier.


To tackle the first question, we estimate an equation of the following form:

$$\ln \left(\frac{C_{i,t+1}}{C_{i,t-1}} \right) = \beta_o + \beta_1 2yr_t + \beta_2 (LTV_{i,t-1} \cdot 2y_t) + \beta_3 LTV_{i,t-1} + \mathbf{BZ}_{i,t} + \lambda_{s(t)} + u_{i,t}. \tag{3}$$

As earlier, the model is estimated using GMM, where the two-year yield is instrumented by monetary policy shocks. The equation is estimated for a subset of the entire sample, i.e., homeowners with mortgages,¹⁴ as housing-related leverage would be a relevant

¹⁴Following the crisis, the homeownership rate fell, and many households also lost their homes to foreclosures. While the CEX data reflect the distribution of homeownership and leverage in the sample, we are unable to identify households that foreclosed. As such, the regressions are based only on homeowners with mortgages.

Figure 3. Effect of Monetary Policy on Consumption Growth at Different LTV Levels



Note: This graph shows the response of consumption to a 10 basis point increase in the two-year yield at different levels of household indebtedness. The x-axis denotes LTV ratio and the y-axis is the consumption response measured by $\beta_1 + \beta_2 * LTV$ (based on Equation (3)), estimated at different LTV levels and corresponding 90 percent confidence intervals.

parameter for this group. A negative value of β_2 supports the hypothesis that households with higher indebtedness respond more to monetary policy shocks. However, the total effect of monetary policy loosening on consumption growth must be read from $\beta_1 + \beta_2 * LTV$.

The results show that β_2 has a negative sign, in line with the notion of a higher responsiveness of more indebted households. The estimated coefficient is, however, only significant for durable consumption. To understand further whether and how the responsiveness of consumption to monetary policy shocks varies with household indebtedness, we check for the joint significance of β_1 and β_2 along the spectrum of possible values for indebtedness (Figure 3).

The confidence intervals widen at higher loan-to-value (LTV) levels. Furthermore, in the case of durable consumption, the overall impact of monetary policy is found to be significant only at LTV levels higher than 0.5, suggesting that monetary policy may be effective only beyond a certain threshold. These results indicate that the response of consumption growth to monetary policy may potentially be non-linear.

As discussed earlier, looser monetary policy would be expected to strengthen balance sheets and reduce indebtedness by boosting

house prices and reducing the net present value of mortgage payments. In turn, these effects should relax credit constraints and favor higher consumption. However, when indebtedness is especially high, the marginal improvement in balance sheets may not be sufficient to restore access to credit or allow for debt refinancing. Moreover, as discussed in Alpanda and Zubairy (2019), high levels of debt may dampen the effectiveness of monetary policy because “highly indebted households may be less willing, or less able, to borrow further in response to a rate cut, especially during recessionary periods when agents are facing higher job insecurity and income uncertainty.” After a shock, households may need to rebuild wealth and increase precautionary savings (Carroll and Kimball 1996, Mian and Sufi 2014). A more specific channel refers to the mechanism by which underwater households may not invest in their homes in response to a monetary easing (Melzer 2017). We refer to the potentially dampening effect of high debt through these various channels as the debt-overhang hypothesis.


To study whether consumption growth responds non-linearly to household indebtedness, we (i) estimate an alternative specification with a quadratic term in LTV and (ii) estimate a threshold regression.

The quadratic specification is as follows:

$$\begin{aligned} \ln \left(\frac{C_{i,t+1}}{C_{i,t-1}} \right) = & \beta_o + \beta_1 2yr_t + \beta_2 (LTV_{i,t-1} \cdot 2y_t) \\ & + \beta_3 (LTV_{i,t-1}^2 \cdot 2y_t) + \beta_4 LTV_{i,t-1} + \beta_5 LTV_{i,t-1}^2 \\ & + \mathbf{BZ}_{i,t} + \lambda_{s(t)} + u_{i,t}. \end{aligned} \quad (4)$$

We focus on the total marginal effect ($\beta_1 + \beta_2 LTV + \beta_3 LTV^2$) to assess the marginal effect of monetary policy on consumption along the spectrum of LTV values. The results (Figure 4) are largely in line with our earlier findings from Figure 3. As before, for non-durable consumption, the marginal effect increases even at higher levels of LTV. However, for durable consumption the negative effect is only significant for LTVs between 0.5 and 1.4—i.e., the effect is not significant at very low nor very high levels of indebtedness.

Figure 4. Effect of Monetary Policy on Consumption Growth at Different LTV Levels Using a Quadratic Specification



Note: This graph shows the response of consumption to a 10 basis point increase in the two-year yield at different levels of household indebtedness. The x-axis denotes LTV ratio and the y-axis is the consumption response measured by $\widehat{\beta}_1 + \widehat{\beta}_2 * LTV + \widehat{\beta}_3 LTV^2$ (based on Equation (4)), estimated at different LTV levels and corresponding 90 percent confidence intervals.

Second, we estimate a threshold regression of the following form:

$$\ln \left(\frac{C_{i,t+1}}{C_{i,t-1}} \right) = \beta_o + \beta_1 2yr_t + \beta_2 (I_{LTV < 0.9} \cdot 2yr_t) + \beta_3 I_{LTV < 0.9} + \mathbf{BZ}_{i,t} + \lambda_{s(t)} + u_{i,t}, \quad (5)$$

where $I_{LTV < 0.9}$ is an indicator function that takes a value of 1 when indebtedness is less than the 90th percentile (which corresponds to an LTV value of 0.99) over the sample period 1996:Q1–2014:Q4. Therefore, a significant value of β_2 implies that transmission is different across households with high and low indebtedness. For LTVs lower than the 90th percentile, the marginal effect of monetary policy on consumption is $(\beta_1 + \beta_2)$, while it is β_1 for LTVs in the top 10 percentile. We find that higher indebtedness increases the responsiveness to monetary policy shocks for non-durable consumption over the full and pre-crisis samples (Table 6, columns 3 and 5, respectively). Thus, the effects of indebtedness appear to be non-linear. The results for durable consumption are comparable over the pre-crisis sample (Table 6, column 6), and have the expected

Table 6. Impact of Monetary Policy on Consumption: The Role of Indebtedness

Variables	Non-durables (1)	Durables (2)	Non-durables (3)	Durables (4)	Non-durables (5)	Durables (6)
Two-Year Yield	-22.41*** (7.19)	5.66 (6.02)	-8.73*** (2.81)	-27.33 (22.05)	-18.93*** (4.30)	-41.46* (23.41)
Family Size	-0.15 (0.14)	-0.16 (1.02)	-0.07 (0.12)	-0.23 (1.01)	-0.10 (0.14)	-1.36 (1.21)
College Education	0.77** (0.35)	0.81 (2.71)	0.97*** (0.32)	1.87 (2.55)	0.90** (0.39)	2.45 (3.06)
White	0.89* (0.53)	0.31 (4.23)	1.19*** (0.45)	-0.63 (4.17)	0.72 (0.56)	-3.49 (5.23)
Married	0.57 (0.42)	0.16 (3.35)	0.41 (0.38)	-1.26 (3.28)	0.39 (0.48)	-0.93 (4.01)
Reference Age	0.01 (0.02)	-0.15 (0.14)	0.03** (0.01)	-0.10 (0.11)	0.03 (0.02)	-0.13 (0.14)
LTV	29.48 (28.66)	69.79** (29.68)				
LTV*2-yr Yield	-10.67 (10.11)	-21.43** (10.04)				
I.(LTV < 0.9)			-6.69*** (1.78)	-6.43 (14.65)	-40.22*** (8.66)	-78.48 (48.50)
I.(LTV < 0.9)*2-yr Yield			2.44*** (0.67)	5.02 (5.15)	9.74*** (2.10)	19.79* (11.31)
Observations	73,184	36,866	73,307	37,225	51,272	26,713
No. of Households	37,576	24,351	37,644	24,517	26,075	17,346
Hansen	Exactly Identified	0.729	Exactly Identified	0.222	Exactly Identified	0.515
Sample	Full	Full	Full	Full	Pre-crisis	Pre-crisis

Note: GMM estimation, 1996:Q1–2014:Q4. Dependent variable is two-quarter-ahead consumption growth. In the first-stage regression, two-year yield is instrumented by monetary policy shocks. All regressions include a constant, and quarter (seasonal) effects. Clustered standard errors (by households) are reported in parentheses. ***, **, and * indicate statistical significance at the 1, 5, and 10 percent levels, respectively.

sign, though they are not significant for the full sample (Table 6, column 4).

We further explore the responsiveness of consumption at other thresholds, namely at 70th (0.75), 80th (0.86), and 95th (1.26) percentiles (LTV values). The results corroborate the earlier findings: the response to monetary policy shocks increases with indebtedness, but there is no evidence that it diminishes at high levels (Table 7). This is particularly evident for durable consumption, which shows a monotonically increasing coefficient on indebtedness as the threshold is raised. These findings contrast with literature that finds a diminished response to monetary policy at high levels of indebtedness (Beraja et al. 2019).

In summary, our results suggest that more indebted households respond more to monetary policy impulses. However, we do not find evidence of a debt overhang effect—that is, of a weakened response at very high levels of indebtedness.

Therefore, for indebtedness to explain the decrease in the average response of household consumption to monetary policy shocks, the overall distribution of indebtedness must have shifted leftward post-crisis, toward less indebted households. However, if anything, the distribution of indebtedness shifted to the right,¹⁵ though its mean declined somewhat, as shown in Figure 5. Using estimated coefficients from Equation (3) (Figure 3), the responsiveness of non-durable and durable consumption to a 10 basis point rise in the two-year yield is found to increase by 2 and 4 basis points, respectively, due to the shift in the distribution of household indebtedness post-crisis. The proportion of households in the top 10 percentile of LTV distribution grew from 5 percent before the crisis to 8 percent in the post-crisis period. According to Equation (4), this implies a 3 and 6 basis point increase in the responsiveness of non-durable and durable consumption, respectively, to a 10 basis point hike in the two-year yield. Thus, both specifications indicate that changes


¹⁵While Guren et al. (2018) do not study the responsiveness of consumption to monetary policy shocks, to the extent that monetary policy has an impact on house prices, they find that the housing wealth effect is not sensitive to changes in the distribution of LTV because a rightward shift in the LTV distribution not only increases the number of highly sensitive constrained agents but also the number of underwater agents whose consumption is insensitive to house prices.

Table 7. Response of Consumption Growth to a 10 Basis Point Increase in the Two-Year Yield at Different LTV Thresholds¹ (in percentage points)

	LTV < 70	LTV > 70	LTV < 80	LTV > 80	LTV < 95	LTV > 95
<i>Full Sample</i>						
Non-durables	-0.63*** (0.004)	-0.72*** (0.002)	-0.63*** (0.004)	-0.70*** (0.002)	-0.63*** (0.004)	-1.09*** (0.002)
Durables	-2.19 (0.201)	-2.35 (0.203)	-2.22 (0.198)	-2.55 (0.185)	-2.24 (0.193)	-3.49 (0.204)
<i>Pre-crisis</i>						
Non-durables	-0.91*** (0.000)	-1.19*** (0.000)	-0.92*** (0.000)	-1.30*** (0.000)	-0.94*** (0.000)	-2.99*** (0.000)
Durables	-2.21* (0.085)	-2.50 (0.102)	-2.19* (0.087)	-2.93* (0.083)	-2.14* (0.085)	-6.77* (0.052)

¹Figures in parentheses are P-values.

Figure 5. Density of Loan-to-Value




in the LTV distribution have per se contributed to a higher responsiveness of consumption to monetary policy in the post-crisis period. We must therefore look elsewhere to seek a plausible explanation for the drop in monetary policy effectiveness relative to consumption.

While the distribution of household indebtedness in CEX does not appear to have shifted significantly post-crisis, the survey also allows us to study the evolution of homeowners’ interest rate exposure. Studies using household-level data in countries where variable rate mortgages are the norm find that the consumption responsiveness to monetary policy increases as the interest rate exposure increases (Holm, Pascal, and Tischbirek 2021). Here, we examine the evolution of the share of non-FRMs, i.e., mortgages that do not carry a fixed interest rate over the term of the loan. A higher share of non-FRMs implies a larger share of households with interest rate exposure.

The data show that while FRMs account for the majority of mortgages in the United States, the share of non-FRMs dropped significantly at the beginning of the crisis in 2007, and it has stayed low since then (Figure 6). Such a sharp decline may have been driven by a lower risk appetite of households at the onset of the crisis. The lower interest rates following the crisis further encouraged households to shift towards FRMs and fix low interest payments for the entire term of the mortgage (Wilson 2016).

**Figure 6. Non-Fixed-Rate Mortgages
as a Share of All Mortgages**



For households with FRMs, monetary transmission operates primarily through the balance sheet channel, whereby monetary policy can have an impact on households’ net wealth. By contrast, for households with non-FRMs, monetary transmission operates through both the balance sheet channel and the debt-service channel, due to the interest rate exposure. Auclert (2019) notes that the interest rate exposure is a key transmission channel of monetary policy to consumption.

We find that while sign of the estimated coefficients is negative (consistent with stronger monetary policy transmission for households with non-FRMs), the effects are not statistically significant. Specifically, the coefficients on the interaction between the non-FRM dummy and two-year yield for both non-durable and durable goods are statistically insignificant (Table 8). We therefore fail to find convincing evidence that the reduced interest rate exposure of households offers a possible explanation for the post-crisis weakening of the monetary policy transmission to consumption.

5. Does Household Liquidity Matter?

We proceed in much the same way as in the earlier section. We ask whether the liquidity position of households affects their

**Table 8. Impact of Monetary Policy on Consumption:
The Role of Interest Exposure**

Variables	Non-durables (1)	Durables (2)
2-yr Yield	−32.97 (35.97)	−18.31 (15.20)
Non-FRM*2-yr Yield	−15.68 (85.60)	−4.46 (7.49)
Family Size	−0.19 (0.56)	−0.19 (1.01)
College Education	0.74 (0.60)	2.40 (2.55)
White	1.25** (0.55)	−0.35 (4.17)
Married	0.91 (1.93)	−1.23 (3.28)
Reference Age	0.02 (0.03)	−0.09 (0.11)
Non-FRM	59.73 (323.42)	19.10 (28.67)
Observations	73,331	36,936
No. of Households	37,656	24,407
Hansen	Exactly Identified	0.134

Note: GMM estimation, 1996:Q1–2014:Q4. Dependent variable is two-quarter-ahead consumption growth. In the first-stage regression, two-year yield is instrumented by monetary policy shocks. All regressions include a constant, and quarter (seasonal) effects. Clustered standard errors (by households) are reported in parentheses. ***, **, and * indicate statistical significance at the 1, 5, and 10 percent levels, respectively.

consumption response to monetary policy impulses. We further ask whether there are non-linearities, and whether the change in the distribution of household liquidity post-crisis might help explain the lower monetary policy impact on consumption detected earlier in this paper.

We begin by estimating the following equation:

$$\ln \left(\frac{C_{i,t+1}}{C_{i,t-1}} \right) = \beta_o + \beta_1 2yr_t + \beta_2 (LIQ_{t-1} \cdot 2yr_t) + \beta_3 LIQ_{t-1} + \mathbf{BZ}_{i,t} + \lambda_{s(t)} + u_{i,t}. \quad (6)$$

Again, this equation is estimated only for households with mortgages. In this specification, a positive value of β_2 supports the hypothesis that households with low liquidity respond more to monetary policy shocks.

Estimates of β_2 are however found to be insignificant (Table 9, columns 1 and 2). To investigate the issue further, we examine whether the responsiveness of consumption to monetary policy shocks varies with liquidity levels. For this purpose, we check for the joint significance of β_1 and β_2 along the spectrum of liquidity values (Figure 7.1). The results show that the responsiveness of non-durable consumption is only significant at relatively low liquidity values (with liquid-assets-to-monthly-income ratios of up to around 1). As in the case of indebtedness, these results are robust to using an alternative specification that is quadratic in the liquidity term (Figure 7.2).

We explore the possibility that only households with liquidity below a certain threshold respond more to interest rate shocks in a non-linear setting. Specifically, we consider the following threshold regressions:

$$\ln \left(\frac{C_{i,t+1}}{C_{i,t-1}} \right) = \beta_o + \beta_1 2yr_t + \beta_2 (I_{LIQ>.25} \cdot 2yr_t) + \beta_3 I_{LIQ>.25} + \mathbf{BZ}_{i,t} + \lambda_{s(t)} + u_{i,t}, \quad (7)$$

where $I_{LIQ>.25}$ is an indicator function that takes a value of 1 when a household's ratio of liquid assets to income is greater than the 25th percentile (which corresponds to a ratio of liquid asset to monthly income of 0.01) over the sample period 1996:Q1–2014:Q4.


The results indicate that non-durable consumption responds most strongly when households are liquidity constrained. We find qualitatively similar, but not statistically significant, results for

Table 9. Impact of Monetary Policy on Consumption: The Role of Liquidity

Variables	Non-durables (1)	Durables (2)	Non-durables (3)	Durables (4)	Non-durables (5)	Durables (6)
Two-Year Yield	-25.58** (10.39)	-2.51 (9.36)	-23.41*** (8.58)	-14.98 (18.26)	-24.92*** (9.27)	-18.62 (21.55)
Liq*2-yr Yield	7.39 (10.56)	-10.63 (28.88)				
Family Size	-0.12 (0.16)	0.02 (1.17)	0.01 (0.14)	0.16 (1.15)	-0.06 (0.14)	-0.18 (1.15)
College Education	1.56** (0.73)	1.52 (3.45)	0.88** (0.38)	3.21 (2.92)	1.08*** (0.38)	4.17 (2.90)
White	0.80 (0.89)	2.87 (5.18)	0.90 (0.55)	1.11 (4.84)	0.91 (0.56)	1.61 (4.84)
Married	0.56 (0.55)	0.54 (3.82)	0.46 (0.45)	-1.19 (3.69)	0.54 (0.46)	-0.45 (3.70)
Reference Age	0.03 (0.02)	-0.25* (0.15)	0.02* (0.01)	-0.12 (0.13)	0.03* (0.02)	-0.12 (0.13)
Liquid Assets/Income	-23.82 (35.52)	39.01 (100.26)				
I.(liq > 0.25)			-4.00* (2.25)	4.28 (9.04)		
I.(liq > 0.25)*2-yr Yield			2.27*** (0.75)	1.21 (2.61)		
I.(liq > 0.10)					-9.16** (4.55)	-18.45 (19.10)
I.(liq > 0.10)*2-yr Yield					3.93*** (1.51)	5.36 (5.81)
Observations	52,344	28,891	52,345	28,892	52,345	28,892
No. of Households	26,885	18,821	26,885	18,821	26,885	18,821
Hansen		0.129		0.0455		0.0439


Note: GMM estimation, 1996:Q1–2014:Q4. Dependent variable is two-quarter-ahead consumption growth. In the first-stage regression, two-year yield is instrumented by monetary policy shocks. All regressions include a constant, and quarter (seasonal) effects. Clustered standard errors (by households) are reported in parentheses. ***, **, and * indicate statistical significance at the 1, 5, and 10 percent levels, respectively.

Figure 7.1. Effect of Monetary Policy on Consumption Growth at Different Liquidity Levels



Note: This graph shows the response of consumption to a 10 basis point increase in two-year yield at different levels of household liquidity. The x-axis denotes liquid-assets-to-income ratio and the y-axis is the consumption response, measured by $(\widehat{\beta}_1 + \widehat{\beta}_2 * LIQ)$ (based on Equation (5)), estimated at different liquidity levels and corresponding 90 percent confidence intervals.

Figure 7.2. Effect of Monetary Policy on Consumption Growth at Different Liquidity Levels Using a Quadratic Specification



Note: This graph shows the response of consumption to a 10 basis point increase in two-year yield at different levels of household liquidity. The x-axis denotes liquid-assets-to-income ratio and the y-axis is the consumption response, measured by $(\widehat{\beta}_1 + \widehat{\beta}_2 * LIQ + \widehat{\beta}_3 * LIQ^2)$, estimated at different liquidity levels and corresponding 90 percent confidence intervals.

durable consumption (Table 9, columns 4 and 6).¹⁶ Table 10 offers an interpretation of the results, listing the extent of the consumption response to a 10 basis point surprise hike in the two-year interest rate. The response of non-durable consumption increases monotonically as liquidity is lowered from the 20th to the 10th and 5th percentiles. In the first case, consumption of non-durables rises by 2.3 percentage points, while in the last it increases by 2.5 percentage points—not an innocuous difference. Appendix C shows that these results hold if we broaden the definition of liquid assets to include securities such as stocks, mutual funds, private bonds, government bonds, or Treasury notes.

Overall, our results provide some support for the findings of Kaplan and Violante (2014) that non-durable consumption of wealthy hand-to-mouth households (namely those with homeownership but limited liquid assets) responds more strongly to interest rate changes.

Lastly, we ask whether the change in the distribution of liquidity from pre- to post-crisis times might help explain the decline in monetary policy effects on consumption. For liquidity to be relevant, the distribution should have moved rightward, toward a lower share of liquidity-constrained and highly responsive households.

However, the distribution of liquidity has hardly changed over time, or, if anything, has shifted to the left (Figure 8). Based on coefficient estimates from Equation (5), the responsiveness of non-durable consumption is found to marginally strengthen after the crisis due to the observed shift in the liquidity distribution (a 10 basis point increase in two-year yields leads to an additional 0.1 basis point decline in non-durable consumption in the post-crisis period).

The share of households in the lower 25th percentile of liquidity distribution rose from 24 percent pre-crisis to 28 percent after the crisis, which according to the estimates from Equation (6) should also strengthen the responsiveness of non-durable consumption by 0.1 basis point (to a 10 basis point increase in two-year yield).


¹⁶As in the case of leverage, we also estimated Equation (6) for durable consumption for the pre-crisis period. Liquidity continues to not matter for transmission of monetary policy to durable consumption even in the pre-crisis period.

Table 10. Response of Consumption Growth to a 10 Basis Point Increase in the Two-Year Yield at Different Liquid-Asset-to-Income Thresholds (in percentage points)

	LIQ < 25	LIQ > 25	LIQ < 20	LIQ > 20	LIQ < 10	LIQ > 10
Non-durables	-2.34*** (0.006)	-2.11*** (0.007)	-2.39*** (0.006)	-2.11*** (0.007)	-2.49*** (0.007)	-2.10*** (0.007)
Durables	-1.50 (0.412)	-1.38 (0.399)	-1.72 (0.364)	-1.36 (0.408)	-1.86 (0.387)	-1.33 (0.419)

Note: Figures in parentheses are P-values.

Figure 8. Density of Liquid-Asset-to-Income



Hence, changes in liquidity cannot explain the weakened response of consumption observed after the crisis either. The widespread concern that a deterioration of household balance sheets after the crisis hampered monetary policy effectiveness thus does not seem to hold. The explanation for the lower effectiveness of monetary policy must therefore lie elsewhere, such as in the higher degree of economic uncertainty brought about by the crisis. Recent studies (Aastveit, Natvik, and Sola 2017, Castelnuevo and Pellegrino 2018) find that U.S. monetary policy shocks affect economic activity less when uncertainty is high, in line with “real-option” effects from theory (e.g., Bloom 2009). While not reported here, we explored whether the higher uncertainty in the post-crisis period accounts for the lower effectiveness of monetary policy, by interacting monetary policy shocks in our estimation with the index of economic policy uncertainty developed by Baker, Bloom, and Davis (2016). While this preliminary investigation provides only suggestive evidence for such effects, further research on this issue seems worthwhile.

6. Conclusion

We find that the average responsiveness of U.S. household consumption to well-identified monetary policy shocks has declined since the

global financial crisis. However, this result cannot be explained by higher indebtedness or lower liquidity levels. Households with higher debt levels are the most responsive to monetary policy—which contrasts with the findings in Beraja et al. (2019)—and the share of these households in the population grew. We also fail to find convincing evidence that the reduced interest rate exposure of households explains the post-crisis weakening of the monetary policy transmission to consumption. Similarly, while households with lower share of liquid assets are more responsive to monetary policy—which supports Kaplan and Violante (2014)’s finding that homeowners with limited liquid assets respond strongly to interest rate changes—their share in the population somewhat grew post-crisis. Therefore, the common notion that a deterioration of household balance sheets after the crisis hampered monetary policy effectiveness is not validated in the data.

Nevertheless, household balance sheets do matter for the strength of monetary policy transmission, and our results underscore the notion that monetary policymakers need to pay close attention to them. Moreover, given the presence of non-linearities (the responsiveness of more indebted households rises non-linearly with indebtedness), monitoring the distribution of household balance sheet characteristics is important.

Appendix A. Data

The Consumer Expenditure Survey (CEX) is a survey conducted by the Census Bureau and is primarily used by the Bureau of Labor Statistics to determine the weights assigned to different goods and services in calculating the consumer price index (CPI). The CEX is a rotating panel survey, and each household is interviewed once per every 3 months for, at most, 15 consecutive months. In addition, the survey sample is designed to be representative of the U.S. civilian non-institutional population.

A.1 Data Cleanup

We take several steps to clean up the raw CEX data. We drop observations in which the CEX records negative consumption for households. We also drop observations for households with more

than one consumption unit and households with less than four interview observations. This cleanup results in roughly 5,000 quarterly household observations, of which 74 percent are homeowners and 45 percent are homeowners with outstanding mortgage balance. On average, households spend \$4,320 on non-durable goods and \$1,048 on durable goods. Some summary statistics for housing tenure and consumption for the data are shown in Table A.1.

One peculiar feature of the CEX survey is that the interview quarter and the consumption quarters may not align perfectly. Each time a household is interviewed, they are asked about their consumption expenditures over the three months prior to the month of interview. Since households may be interviewed during any month within a given quarter, the interview quarter does not necessarily correspond with the months for which the consumption data are acquired. We make the appropriate adjustments to the consumption data so that they align with their respective calendar quarters.

A.2 Definitions of Consumption Variables

Following Aladangady (2014), non-durable consumption consists of food, alcohol, tobacco, housing operations, utilities, gasoline, public transportation, personal care, reading, entertainment, apparel, healthcare, and education expenses. Durable consumption consists of expenditure on cars (new and used), furniture, and equipment. Table A.1 outlines the CEX variables used to construct non-durable and durable consumption variables.

Details of the CEX variables used in constructing non-durable and durable consumption variables are mentioned in Table A.2.

A.3 Leverage and Liquidity

Most of the household balance sheet data are only available in the fifth interview, while mortgage information is asked in every interview. Leverage is proxied by the ratio of mortgage balance to the reported house value. We aggregate the mortgage balances reported on all the properties owned by the household. The CEX variables used for constructing this are QBLNCM1X or QBLNCM2X, which report the household's mortgage balance at the beginning of the month, three months prior to the interview, or two months prior to

Table A.1. Housing Tenure and Real Consumption

	Share	Average Level		Average Growth	
		Non-durable	Durable	Non-durable	Durable
Homeowners	74	4,803	1,204	0.14	-3.44
with Mortgage	45	5,286	1,397	0.02	-4.02
without Mortgage	29	4,044	900	0.32	-2.25
Renters	26	2,935	601	-0.23	-4.26
All	100	4,320	1,048	0.04	-3.60
Source: CEX and authors' calculations.					
Note: Consumption variables are in constant dollars (2000:Q1 = 100) and winsorized at 1 percent of each tail.					

Table A.2. Definitions of Key Consumption Variables

Variables	Details	CEX Name
<i>C</i>	Total Expenditure	TOTEXP
	Non-durable Expenditure	
	Food	FOOD
	Alcohol	ALCBEV
	Tobacco	TOBACC
	Housing Operations	HOUSOP
	Utilities	UTIL
	Gasoline	GASMO
	Public Transportation	PUBTRA
	Personal Care	PERSCA
	Reading	READ
	Entertainment	ENTERT
	Apparel	APPAR
	Healthcare	HEALTH
	Educational Expenses	EDUCA
	Durable Expenditure	
	Cars and Trucks, New	CARTKN
	Cars and Trucks, Used	CARTKU
	Other Vehicles	OTHVEH
	Furnishing and Equipment	HOUSEQ

the interview, respectively. Our choice over which of the two variables to use depends on which month corresponds to the first month in the consumption quarter. If a household refinances its mortgage on a property, we adjust the household’s mortgage balance such that the mortgage balances before and after refinancing are not double-counted. For property value we use PROPVALX. We construct a house price index using this variable and it matches well with the Case-Shiller Home Price Index, particularly the boom-bust in the house prices, although it is not shown here.

Liquid assets include the total balance a household has in their checking and savings accounts. From 2013 onwards, liquid assets also include money market accounts and certificates of deposits. The CEX variables used in constructing the liquid assets variable are LIQUIDX for the period covering 2013–14 and CKBKACTX + SAVACCTX for 1994–2012. Unlike balance sheet variables, income

is reported in both the second and the fifth interview. We use the imputed after-tax income, FINCATXM, from 2004 onwards. For the prior years, we use the reported after-tax income, FINCATAX, and replace invalid missing entries with imputed income data. Table A.3 shows the correlation matrix among key variables.

A.4 Cohorts and Control Variables

We construct the synthetic cohorts using housing tenure (CUTENURE) and the household head's birth year, which is determined by the interview date and the household head's age (AGE_REF) at the time of the interview. The control variables used in the panel analysis include race (REF_RACE), education (EDUC_REF), age (AGE_REF), family size (FAM_SIZE), and marital status (MARITAL1).

Appendix B. Synthetic Cohort Panel Data

B.1 Construction of Synthetic Cohort Panel

To measure the responsiveness of households' consumption to monetary policy over time, we need a panel data, although the CEX is designed as repeated cross-section data (Appendix A). Therefore, we construct a synthetic panel, as in Attanasio and Davis (1996), Narita and Narita (2011), and Cloyne, Ferreira, and Surico (2018). We construct synthetic cohorts based on the arguably time-invariant household characteristics, which are the birth year and housing tenure of the household head. That is, we construct a panel data set of each representative consumer unit (CU) with one of the combination of these characteristics.

The birth cohorts are defined by a five-year band. The oldest cohort consists of people who were born between January 1910 and December 1914. We focus on household heads of age 25 to 75. The housing status is categorized into three levels: owners without mortgage, owners with positive mortgage balance, and renters. This procedure yields an unbalanced panel of 42 synthetic cohorts with a minimum of 20 CUs in each of them.

Our choice of a small set of characteristics is driven by the objectives of avoiding having few CUs in some synthetic cohorts, and

Table A.3. Correlation Matrix

	Non-durable Consumption	Durable Consumption	LTV	Liquidity	Family Size	College Education	Ethnicity (White = 1)	Marital Status	Reference Age
Non-durable Consumption	1								
Durable Consumption	0.045	1							
LTV	-0.0068	-0.0068	1						
Liquidity	0.0098	0.0039	-0.1774	1					
Family Size	0.0001	0.0004	0.2209	-0.1502	1				
College Education	0.0055	0.004	0.0839	0.0879	-0.0121	1			
Ethnicity (White = 1)	0.0056	0.0028	-0.0481	0.0802	-0.0351	0.0374	1		
Marital Status	0.0055	0.0027	0.0821	-0.0056	0.4539	0.1051	0.1163	1	
Reference Age	0.0078	0.003	-0.4468	0.2449	-0.3801	-0.088	0.0589	-0.0947	1

avoiding short time series. The number of CUs in a synthetic cohort varies across cohorts. This variation in the number of CUs in synthetic cohorts can be problematic. The time-series data of synthetic cohorts with few CUs tend to be much volatile than that of synthetic cohorts with many CUs, because household-specific changes in consumption are not averaged out. This leads to high standard errors for synthetic cohorts with few CUs. Also, if the time-series of consumption and income are too short, estimation may suffer from a small-sample bias.

B.2 Estimation of Cohort-Level Variables

Given the definition of synthetic cohorts, we estimate durable and non-durable consumption paths for each cohort. We consider a reduced-form relationship between cohort-level consumption and individual household-level consumption in the cohort as follows:

$$\log(c_{j,i,t}) = \log(c_{i,t}) + \varepsilon_{j,i,t}, \quad \varepsilon_{j,i,t} \sim i.i.d. (0, \sigma_{i,t}^2),$$

where $c_{j,i,t}$ is consumption level of household j in cohort c at time t , $c_{i,t}$ is cohort-level consumption for cohort i at time t , and $\varepsilon_{j,i,t}$ is a household-specific idiosyncratic shock at time t , which has mean zero and variance $\sigma_{i,t}^2$. That is, we model log of individual consumption as a random draw from a distribution with mean $\log(c_{i,t})$ and variance $\sigma_{i,t}^2$.

In this reduced-form model, the simple average of $\log(c_{j,i,t})$ over households in cohort c at time t is a consistent estimate of $\log(c_{i,t})$ by the law of large numbers. Since the CEX is a random sample from U.S. population, we use the CEX sample weights in taking the average. We interpret the CEX sample weights as the number of off-sample households who are represented by the consumer unit in the sample. Namely, we consider that there are $\omega_{j,i,t}$ households who are similar to household i , and hence whose consumptions are equal to $c_{j,i,t}$. Therefore, our estimate of cohort-level logged consumption is the weighted average of logged consumption expenditures over households in the cohort, using the CEX sample weights. That is,

$$\widehat{\log(c_{i,t})} := \frac{1}{\omega_{i,t}} \sum_{j \in I_{i,t}} \omega_{j,i,t} \log(c_{j,i,t})$$

$$\log(c_{i,t}) := 1,$$

where $I_{i,t}$ is the set of households in cohort i at time t , $\omega_{j,i,t}$ is the CEX sample weights, and $\omega_{i,t} := \sum_{j \in I_{i,t}} \omega_{j,i,t}$.

Appendix C. Using Broader Definition of Liquidity

In this section, we expand our definition of liquid assets to include securities such as stocks, mutual funds, private bonds, government bonds, or Treasury notes (SECESTX in CEX), in addition to funds in the checking, savings and money market accounts, and certificates of deposits (which is the definition of liquid assets used in the main text). For households who have a missing value for the estimated value of securities, we assume that they do not possess any securities. Since only a small fraction of households reports their holding of securities, this assumption is needed to ensure that the sample size is the same as the one in Section 5.

The results in Table C.1 show that broadening the definition of liquid assets to include securities does not significantly alter the results derived from using the narrower definition of liquidity. Our results in this exercise suggest that liquidity-constrained households are more responsive to monetary policy shocks. However, as is the case with using the narrower definition of liquidity, the distribution of the broadly defined liquid assets-to-income ratio appears to have shifted leftwards (Figure C.1) following the crisis. This shift implies a higher responsiveness of consumption to monetary policy shocks.

Appendix D. Using Alternative Monetary Policy Shocks

This appendix uses policy news shocks constructed by Nakamura and Steinsson (2018). Results are largely robust to using alternative monetary policy shock.

Figure C1. Density of Broader Liquid-Assets-to-Income




Table D.1. Impact of Monetary Policy on Consumption

Variables	Non-durables (1)	Durables (2)	Non-durables (3)	Durables (4)
Two-Year Yield	−10.59*** (3.68)	−39.80** (15.86)	−43.89*** (8.89)	−18.80 (11.50)
Family Size	−0.15 (0.10)	0.31 (0.88)	−0.14 (0.10)	0.32 (0.87)
College Education	0.90*** (0.26)	1.16 (2.24)	0.99*** (0.27)	1.15 (2.24)
White	0.42 (0.32)	2.83 (3.35)	0.36 (0.33)	2.76 (3.35)
Married	0.55** (0.27)	−0.06 (2.57)	0.52* (0.28)	−0.14 (2.57)
Reference Age	0.02*** (0.01)	0.05 (0.08)	0.03*** (0.01)	0.05 (0.08)
GFC			−17.99*** (3.85)	−1.95 (7.33)
GFC*2-yr Yield			9.26*** (2.56)	12.95 (9.96)
Observations	130,396	48,798	130,396	48,798
No. of Households	67,698	33,555	67,698	33,555
Hansen	Exactly Identified	0.137	Exactly Identified	0.0812

Note: GMM estimation, 1996:Q1–2014:Q4. Dependent variable is two-quarter-ahead consumption growth. In the first-stage regression, two-year yield is instrumented by monetary policy shocks. All regressions include a constant, and quarter (seasonal) effects. Clustered standard errors (by households) are reported in parentheses. ***, **, and * indicate statistical significance at the 1, 5, and 10 percent levels, respectively.

References

- Aastveit, K. A., G. J. Natvik, and S. Sola. 2017. "Economic Uncertainty and the Influence of Monetary Policy." *Journal of International Money and Finance* 76 (September): 50–67.
- Aladangady, A. 2014. "Homeowner Balance Sheets and Monetary Policy." Finance and Economics Discussion Series No. 2014–98, Board of Governors of the Federal Reserve System.
- . 2017. "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata." *American Economic Review* 107 (11): 3415–46.
- Alpanda, S., and S. Zubairy. 2019. "Household Debt Overhang and Transmission of Monetary Policy." *Journal of Money, Credit and Banking* 51 (5): 1265–1307.
- Alter, A., A. X. Feng, and N. Valckx. 2018. "Understanding the Macro-financial Effects of Household Debt: A Global Perspective." IMF Working Paper No. 18/76.
- Attanasio, O. P., and S. Davis. 1996. "Relative Wage Movements and the Distribution of Consumption." *Journal of Political Economy* 104 (6): 1227–62.
- Auclert, A. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109 (6): 2333–67.
- Baker, S. R., N. Bloom, and S. J. Davis. 2016. "Measuring Economic Policy Uncertainty." *Quarterly Journal of Economics* 131 (4): 1593–1636.
- Beraja, M., A. Fuster, E. Hurst, and J. Vavra. 2019. "Regional Heterogeneity and the Refinancing Channel of Monetary Policy." *Quarterly Journal of Economics* 134 (1): 109–83.
- Bernanke, B. S. 2018. "The Real Effects of the Financial Crisis." *Brookings Papers on Economic Activity*, BPEA Conference Drafts, September 13–14.
- Bernanke, B. S., and K. N. Kuttner. 2005. "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *Journal of Finance* 60 (3): 1221–57.
- Bloom, N. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 (3): 623–85.
- Caballero, R. J. 1993. "Durable Goods: An Explanation for Their Slow Adjustment." *Journal of Political Economy* 101 (2): 351–84.

- Carroll, C. D., and M. S. Kimball. 1996. "On the Concavity of the Consumption Function." *Econometrica* 64 (4): 981–92.
- Castelnuovo, E., and G. Pellegrino. 2018. "Uncertainty-Dependent Effects of Monetary Policy Shocks: A New-Keynesian Interpretation." *Journal of Economic Dynamics and Control* 93 (August): 277–96.
- Christiano, L. J., and M. Eichenbaum. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45.
- Christiano, L. J., M. Trabandt, and K. Walentin. 2010. "DSGE Models for Monetary Policy Analysis." In *Handbook of Monetary Economics*, Vol. 3, ed. B. M. Friedman and M. Woodford, 285–367. Amsterdam: North-Holland.
- Cloyne, J., C. Ferreira, and P. Surico. 2018. "Monetary Policy when Households Have Debt: New Evidence on the Transmission Mechanism." Documentos de Trabajo No. 1813, Bank of Spain.
- de Francisco, E. 2019. "Housing Choices and Their Implications for Consumption Heterogeneity." International Finance Discussion Series No. 1249, Board of Governors of the Federal Reserve System.
- Di Maggio, M., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao. 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review* 107 (11): 3550–88.
- Drehmann, M., M. Juselius, and A. Korinek. 2017. "Accounting for Debt Service: The Painful Legacy of Credit Booms." BIS Working Paper No. 645.
- Dynan, K. 2012. "Is a Household Debt Overhang Holding Back Consumption?" *Brookings Papers on Economic Activity* (Spring): 299–362.
- 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.
- Eggertsson, G. B., N. R. Mehrotra, and J. A. Robbins. 2017. "A Model of Secular Stagnation: Theory and Quantitative Evaluation." NBER Working Paper No. 23093.

- Ferrari, M., J. Kearns, and A. Schrimpf. 2017. "Monetary Policy's Rising FX Impact in the Era of Ultra-low Rates." BIS Working Paper No. 626.
- Flodén, M., M. Kilström, J. Sigurdsson, and R. Vestman. 2021. "Household Debt and Monetary Policy: Revealing the Cash-Flow Channel." *Economic Journal* 131 (636): 1742–71.
- Gilchrist, S., D. López-Salido, and E. Zakrajšek. 2015. "Monetary Policy and Real Borrowing Costs at the Zero Lower Bound." *American Economic Journal: Macroeconomics* 7 (1): 77–109.
- Guerrieri, V., and G. Lorenzoni. 2011. "Credit Crises, Precautionary Savings, and the Liquidity Trap." NBER Working Paper No. 17583.
- Guren, A., A. McKay, E. Nakamura, and J. Steinsson. 2018. "Housing Wealth Effects: The Long View." NBER Working Paper No. 24729.
- 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.
- . 2007. "Market-Based Measures of Monetary Policy Expectations." *Journal of Business and Economic Statistics* 25 (2): 201–12.
- Hanson, S. G., and J. C. Stein. 2015. "Monetary Policy and Long-Term Real Rates." *Journal of Financial Economics* 115 (3): 429–48.
- Hedlund, A., F. Karahan, K. Mitman, and S. Ozkan. 2017. "Monetary Policy, Heterogeneity, and the Housing Channel." 2017 Meeting Papers No. 1610, Society for Economic Dynamics.
- Hofmann, B., and G. Peersman. 2017. "Is There a Debt Service Channel of Monetary Transmission?" *BIS Quarterly Review* (December): 23–37.
- Holm, M. B., P. Paul, and A. Tischbirek. 2021. "The Transmission of Monetary Policy under the Microscope." Working Paper No. 2020-03, Federal Reserve Bank of San Francisco.
- Iacoviello, M., and S. Neri. 2010. "Housing Market Spillovers: Evidence from an Estimated DSGE Model." *American Economic Journal: Macroeconomics* 2 (2): 125–64.
- International Monetary Fund. 2017. "Household Debt and Financial Stability." Chapter 2 in *Global Financial Stability Report*:

- Is Growth at Risk?* (October). Washington, DC: International Monetary Fund.
- Jordà, Ò. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–82.
- Justiniano, A., G. Primiceri, and A. Tambalotti. 2015. "Household Leveraging and Deleveraging." *Review of Economic Dynamics* 18 (January): 3–20.
- Kaplan, G., B. Moll, and G. L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review* 108 (3): 697–743.
- Kaplan, G., and G. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82 (4): 1199–1239.
- . 2018. "Microeconomic Heterogeneity and Macroeconomic Shocks." *Journal of Economic Perspectives* 32 (3): 167–94.
- Kaplan, G., G. Violante, and J. Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity* (Spring): 77–138.
- Luo, W. 2017. "The Consumer Credit Channel of Monetary Policy." Ph.D. dissertation. Retrieved from Georgetown University Digital Georgetown repository, https://repository.library.georgetown.edu/bitstream/handle/10822/1043836/Luo-georgetown_0076D.13549.pdf?sequence=1.
- Melzer, B. T. 2017. "Mortgage Debt Overhang: Reduced Investment by Homeowners at Risk of Default." *Journal of Finance* 72 (2): 575–612.
- 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. "What Explains the 2007-2009 Drop in Employment." *Econometrica* 82 (6): 2197–2223.
- Mian, A., A. Sufi, and E. Verner. 2017. "Household Debt and Business Cycles Worldwide." *Quarterly Journal of Economics* 132 (4): 1755–1817.
- Nakamura, E., and J. Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics* 133 (3): 1283–1330.
- Narita, F., and M. Narita. 2011. "Risk and Risk-Sharing." In *Essays in Uninsurable Income Risk and Household Behavior*,

- ed. M. Narita (chapter 2). Ph.D. dissertation. Retrieved from the University of Minnesota Digital Conservancy, <http://hdl.handle.net/11299/112724>.
- 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.
- Ramey, V. A. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics*, Vol. 2A, ed. J. B. Taylor and H. Uhlig, 71–162 (chapter 2). Elsevier.
- Romer, C. D., and D. H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94 (4): 1055–84.
- Souleles, N. S. 1999. "The Response of Household Consumption to Income Tax Refunds." *American Economic Review* 89 (4): 947–58.
- Stock, J. H., and M. W. Watson. 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *Economic Journal* 128 (610): 917–48.
- Sufi, A. 2015. "Out of Many, One? Household Debt, Redistribution and Monetary Policy during the Economic Slump." Andrew Crockett Memorial Lecture, Bank for International Settlements, Basel, Switzerland, June 28.
- Thimme, J. 2017. "Intertemporal Substitution in Consumption: A Literature Review." *Journal of Economic Surveys* 31 (1): 226–57.
- Wilson, T. J. 2016. "What the Consumer Expenditure Survey Tells Us about Mortgage Instruments Before and After the Housing Collapse." *Beyond the Numbers: Prices & Spending* (U.S. Bureau of Labor and Statistics) 5 (14). <https://www.bls.gov/opub/btn/volume-5/what-the-consumer-expenditure-survey-tells-us-about-mortgage-instruments-before-and-after-the-housing-collapse.htm>.
- Wong, A. 2015. "Population Aging and the Aggregate Effects of Monetary Policy." <https://ssrn.com/abstract=2462555>.
- Yellen, J. 2016. "Macroeconomic Research After the Crisis." Speech delivered at "The Elusive 'Great' Recovery: Causes and Implications for Future Business Cycle Dynamics" 60th annual economic conference sponsored by the Federal Reserve Bank of Boston, Boston, Massachusetts, October 14. <https://www.federalreserve.gov/newsevents/speech/yellen20161014a.htm>.