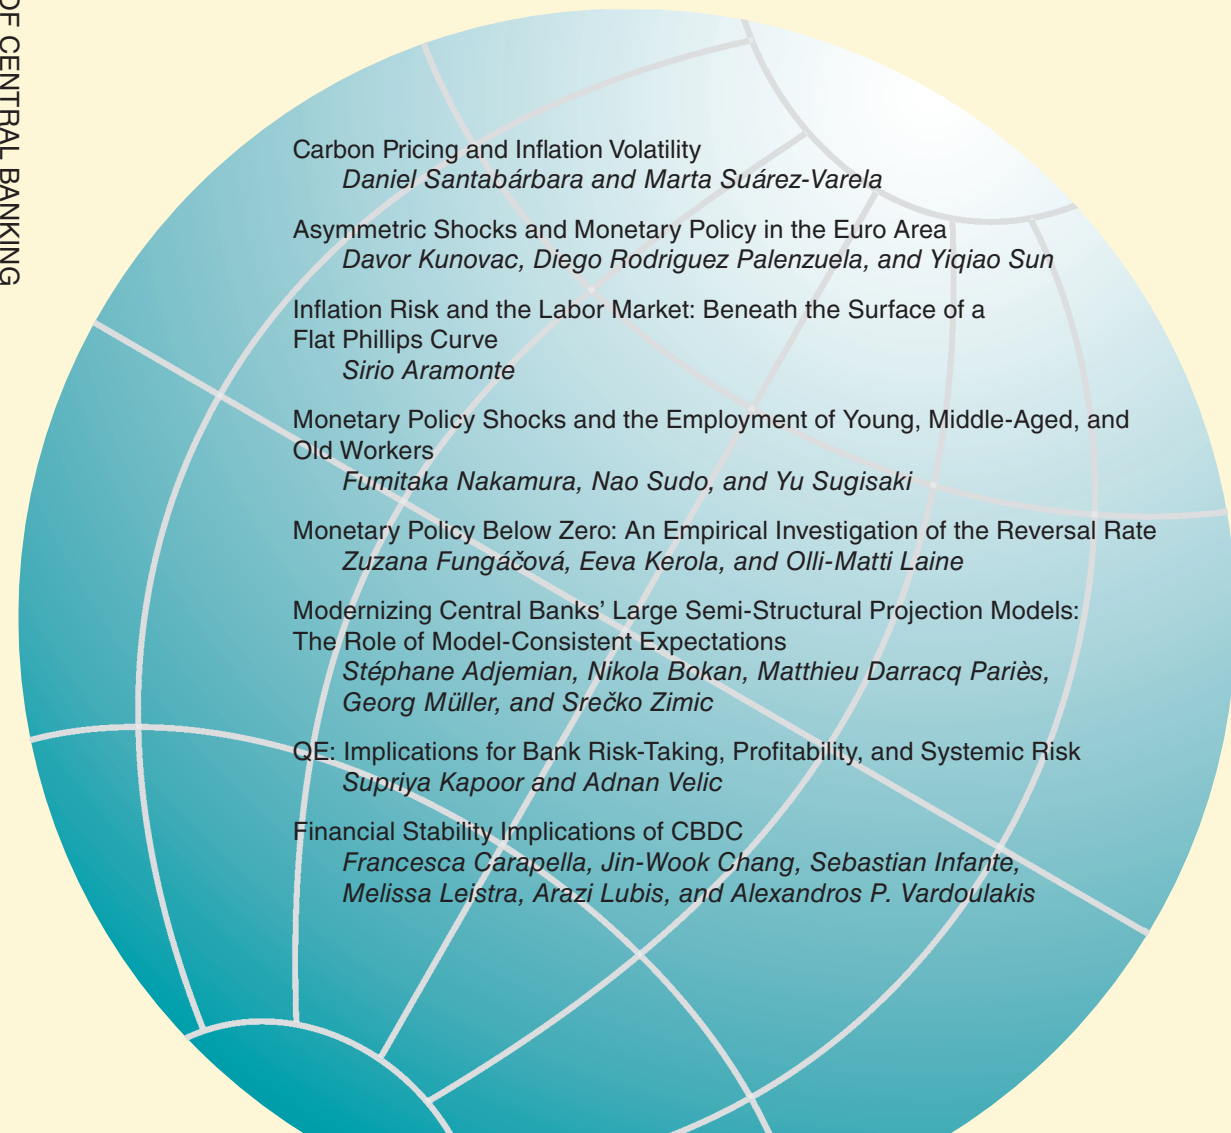




INTERNATIONAL JOURNAL OF CENTRAL BANKING
Volume 22, Number 3
July 2026



INTERNATIONAL JOURNAL OF CENTRAL BANKING



Carbon Pricing and Inflation Volatility
Daniel Santabárbara and Marta Suárez-Varela

Asymmetric Shocks and Monetary Policy in the Euro Area
Davor Kunovac, Diego Rodriguez Palenzuela, and Yiqiao Sun

Inflation Risk and the Labor Market: Beneath the Surface of a Flat Phillips Curve
Sirio Aramonte

Monetary Policy Shocks and the Employment of Young, Middle-Aged, and Old Workers
Fumitaka Nakamura, Nao Sudo, and Yu Sugisaki

Monetary Policy Below Zero: An Empirical Investigation of the Reversal Rate
Zuzana Fungáčová, Eeva Kerola, and Olli-Matti Laine

Modernizing Central Banks' Large Semi-Structural Projection Models: The Role of Model-Consistent Expectations
Stéphane Adjemian, Nikola Bokan, Matthieu Darracq Pariès, Georg Müller, and Srečko Zimic

QE: Implications for Bank Risk-Taking, Profitability, and Systemic Risk
Supriya Kapoor and Adnan Velic

Financial Stability Implications of CBDC
Francesca Carapella, Jin-Wook Chang, Sebastian Infante, Melissa Leistra, Araz Lubis, and Alexandros P. Vardoulakis



Carbon Pricing and Inflation Volatility <i>Daniel Santabárbara and Marta Suárez-Varela</i>	1
Asymmetric Shocks and Monetary Policy in the Euro Area <i>Davor Kunovac, Diego Rodriguez Palenzuela, and Yiqiao Sun</i>	49
Inflation Risk and the Labor Market: Beneath the Surface of a Flat Phillips Curve <i>Sirio Aramonte</i>	119
Monetary Policy Shocks and the Employment of Young, Middle-Aged, and Old Workers <i>Fumitaka Nakamura, Nao Sudo, and Yu Sugisaki</i>	165
Monetary Policy Below Zero: An Empirical Investigation of the Reversal Rate <i>Zuzana Fungáčová, Eeva Kerola, and Olli-Matti Laine</i>	207
Modernizing Central Banks' Large Semi-Structural Projection Models: The Role of Model-Consistent Expectations <i>Stéphane Adjemian, Nikola Bokan, Matthieu Darracq Pariès, Georg Müller, and Srećko Zimic</i>	253
QE: Implications for Bank Risk-Taking, Profitability, and Systemic Risk <i>Supriya Kapoor and Adnan Velic</i>	295
Financial Stability Implications of CBDC <i>Francesca Carapella, Jin-Wook Chang, Sebastian Infante, Melissa Leistra, Arazi Lubis, and Alexandros P. Vardoulakis</i>	337

Copyright © 2026 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**](http://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

Fahad Aldossari, *Saudi Central Bank*
Sergio Nicoletti Altimari, *Banca d'Italia*
Roc Armenter, *Federal Reserve Bank of Philadelphia*
Marco Bardoscia, *Bank of England*
Jan Marc Berk, *The Nederlandsche Bank*
Paul Castillo, *Central Reserve Bank of Peru*
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*
Beverly Hirtle, *Federal Reserve Bank of New York*
Sarah Hunter, *Reserve Bank of Australia*
Mugur Isărescu, *National Bank of Romania*
Esa Jokivuolle, *Bank of Finland*
Edward Knotek, *Federal Reserve Bank of Cleveland*
Anna Kovner, *Federal Reserve Bank of Richmond*
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*

Jae Won Lee, *Bank of Korea*
Carlos Lenz, *Swiss National Bank*
Ye Liu, *People's Bank of China*
Jochen Mankart, *Central Bank of Slovakia*
Karel Mertens, *Federal Reserve Bank of Dallas*
Andre Minella, *Central Bank of Brazil*
Gerard O'Reilly, *Central Bank of Ireland*
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*
Frank Smets, *Bank for International Settlements*
Ulf Söderström, *Sveriges Riksbank*
Salah-Eddine Taleb, *Bank of Algeria*
George Tavlas, *Bank of Greece*
Paula Tkac, *Federal Reserve Bank of Atlanta*
Dobieslaw Tymoczko, *National Bank of Poland*
Hernando Vargas Herrera, *Banco de la República de Colombia*
Shingo Watanabe, *Bank of Japan*
Rafael Wouters, *National Bank of Belgium*
Egon Zakrajšek, *Federal Reserve Bank of Boston*

Editorial Board

Managing Editor

Antoine Martin
Swiss National Bank

Co-editors

Diana Bonfim European Central Bank	Óscar Jordà Federal Reserve Bank of San Francisco	Fernanda Nechio Federal Reserve Bank of San Francisco
Huberto Ennis Federal Reserve Bank of Richmond	Robin L. Lumsdaine Kogod School of Business, American University	Steven Ongena University of Zurich
Carlos Garriga Federal Reserve Bank of St. Louis	Michael McMahon University of Oxford	Enrico Sette European Central Bank
Chao Gu University of Missouri	Emanuel Moench Frankfurt School of Finance & Management	

Associate Editors

James Cloyne University of California, Davis	Sergio Mayordomo Banco de España	Andreas Schrimpf Bank for International Settlements
Emilia Garcia-Appendini University of St. Gallen	Galip Kemal Ozhan International Monetary Fund	Zeynep Senyuz Federal Reserve Board
Sarah Holton European Central Bank	Paolo Pesenti Federal Reserve Bank New York	Nora Traum HEC Montréal
Klodiana Istrefi European Central Bank	Federico Ravenna University of Turin	Kenichi Ueda University of Tokyo
Burcin Kisacikoglu Bilkent University		

Advisory Board

Franklin Allen Imperial College London	Takeo Hoshi University of Tokyo	Carlos Carvalho PUC-Rio
Kristin Forbes MIT Sloan	John Taylor Stanford University	Annette Vissing-Jorgensen Federal Reserve Board

Carbon Pricing and Inflation Volatility*

Daniel Santabárbara and Marta Suárez-Varela
Banco de España

Carbon pricing initiatives, in the form of carbon taxes or emission trading systems (ETSs), have proliferated in the last two decades, providing an opportunity to better understand their economic consequences. This paper assesses the effects of carbon pricing initiatives on inflation volatility. We find evidence of ETS schemes inducing larger volatility in consumer price inflation, while no significant impact is found in the case of carbon taxes. This effect feeds through the energy component, which tends to be more closely related to carbon pricing, while it affects core inflation with some delay, suggesting indirect effects of ETSs on inflation. Finally, we explore several mechanisms contributing to the increase in the volatility of inflation, such as the volatility of underlying carbon prices, the sectoral coverage of these initiatives, the relevance of high-polluting sectors in the productive structure, or the role of free emission allowances in ETS schemes.

JEL Codes: E31, E32, E52, E58, Q48, Q58.

1. Introduction

Not only climate change but also its mitigation policies may affect the business cycle and inflation dynamics. The scientific consensus points to an increasing concern that, without a significant reduction

*The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Banco de España or the Eurosystem. We would like to acknowledge Mario Alloza, Ángel Estrada, Filippo Natoli, Florens Oden-dahl, Javier J. Pérez, Pedro del Río, and the anonymous referees as well as other participants at the Banco de España, European Central Bank, and NGFS seminars, the 1st Workshop of the Eurosystem Research Cluster of Climate Change, the 2nd Workshop of the NGFS Expert Network on Research, and the 2023 Conference of the Association of Environmental and Resource Economists for their valuable comments. All remaining errors are obviously ours. Correspondence to Daniel Santabárbara (daniel.santabarbara@bde.es) and Marta Suárez-Varela (marta.suarez-varela@bde.es).

in greenhouse gas (GHG) emissions, the process of global warming will lead to devastating and potentially irreversible consequences for the planet (Intergovernmental Panel on Climate Change 2023). In light of this evidence, many governments have implemented significant mitigation policies and committed to expanding their ambition. Besides the physical risks derived from climate change, those policies aimed at the transformation to a low-carbon economy could also be associated with significant transition risks,¹ both of which affect macroeconomic outcomes. The realization that climate change and mitigation policies may reshape policy formulations has prompted the need for a better understanding of their consequences for economic policy.²

Among the various policy measures aimed at mitigating climate change, carbon pricing is widely regarded as the most efficient economic instrument for reducing emissions (Hepburn, Stiglitz, and Stern 2020). This approach allows for the internalization of the external costs associated with the production and consumption of GHG-intensive goods, effectively aligning the implicit cost of emissions with their social cost (Pigou 1933). In recent years, the implementation of carbon pricing initiatives has surged, not only in terms of the number of initiatives but also in their geographical spread and the proportion of emissions they cover (World Bank 2021).

Carbon pricing initiatives are designed to increase the prices for carbon emissions, providing economic agents with an incentive to conserve energy and switch to greener sources. Then, they should lead to higher prices of those goods and services intensive in

¹These risks derive from the initiatives to mitigate climate change. For instance, to the extent that many of them lead to an increase in fossil fuel prices in the short term, this could affect the income and credit quality of those households and firms that depend more intensively on this source of energy. Moreover, uncertainty about the public policies and the structural transformation to come could influence economic agents' consumption and investment decisions, or generate financial markets' turbulence episodes. The probability of the realization of these risks and their intensity will depend on how the transition to a low-carbon economy is designed and implemented.

²Both climate change and the transition to a low-carbon economy could affect the delivery of the central bank's price stability mandate. They may hinder the transmission of monetary policy to the extent that these structural processes may induce changes in business cycle and inflation dynamics, affect the equilibrium level of the real interest rate, or generate distortions in the financial system (in particular, in credit allocation).

GHG emissions, pushing up (at least temporarily) inflation. However, incentives from carbon pricing to increase energy efficiency and foster innovation in low-carbon sources of energy, especially in the electricity sector, combined with the diminishing weight of fossil fuels in the consumption basket, act in the opposite direction.³ In this regard, most of the empirical literature so far has focused on the effects of carbon prices on the level of inflation, the interaction between underlying inflationary and disinflationary forces, and, in some cases, on the persistence of inflation (McKibbin, Konradt, and Weder di Mauro 2021; Moessner 2022; Känzig 2023; Konradt and Weder di Mauro 2023). These studies show that carbon pricing initiatives generally produce a transitory effect on the level of headline inflation.

In contrast, inflation volatility has received comparatively less attention.⁴ Inflation volatility is nevertheless crucial for monetary policy, as it may complicate both its communication and implementation. Firstly, sectoral price shocks linked to carbon pricing risk affecting inflation expectations, possibly triggering second-round effects (Batten, Sowerbutts, and Tanaka 2020). Secondly, higher inflation uncertainty has been found to be associated with higher inflation levels (see Cukierman and Meltzer 1986 for a discussion). Along these lines, recent research has shown that inflation volatility can intensify the impact of supply shocks on consumer prices. Specifically, during periods of heightened inflation volatility, the payoffs of swiftly responding to inflation developments may be more significant, thus boosting the incentives for firms to enhance their pricing flexibility (Arndt and Enders 2024). Finally, significant inflation volatility may lead to economic uncertainty, affecting agents' investment and consumption decisions (Judson and Orphanides 1999; Byrne and Davis 2004; Elder 2004), and introducing frictions on markets (Friedman 1977).

³As argued by Schnabel (2022), in order to limit the consequences of climate change, carbon prices should remain elevated in the medium term, while investments and innovation in decarbonized technologies to replace them may need some more time. Therefore, inflationary pressures may be more persistent.

⁴To our knowledge, the impact of carbon pricing on inflation volatility has only been addressed *ex ante* by means of environmental dynamic stochastic general equilibrium (DSGE) (Diluiso et al. 2021), focusing on carbon tax systems.

Importantly, carbon pricing initiatives could be expected to affect the inflation process differently depending on their design. Indeed, several works and reports have already suggested a relationship between certain types of carbon pricing systems and price volatility (see, for example, Diluiso et al. 2021; European Central Bank 2021; Metcalf 2021). This is because the presumed impact on inflation volatility of the two main carbon pricing initiatives, ETS schemes (also known as cap-and-trade schemes) and carbon taxes,⁵ differs. Under an ETS, a cap is first set on pollution, followed by the issuance of a volume of emission allowances compatible with that cap. These allowances can then be traded by companies, determining a market-clearing price that could be potentially rather volatile. Conversely, under carbon tax schemes, policymakers directly set a price on emissions, letting the market determine the amount of pollution consistent with that price. Thus, carbon prices tend to be relatively stable as long as their tax rates and bases do not change frequently. All in all, while carbon taxes deliver relatively stable carbon prices, ETSs have been found to lead to substantial carbon price volatility (Goulder and Schein 2013; Metcalf 2021). Higher volatility of carbon prices could be transmitted into producer and consumer prices—especially of those goods and services related to those sectors under the coverage of carbon pricing—and thus increase inflation volatility. Thus, these initiatives could have only a temporary effect on the level of headline inflation but a more permanent effect on its volatility.

In this paper, we fill this gap in the literature by empirically assessing the effects of the two main types of carbon pricing initiatives (carbon taxes and ETSs) on consumer price inflation volatility. We extend this analysis to evaluate not only the volatility of headline inflation but also its impact on energy and core components, providing a more comprehensive understanding of the channels through which carbon pricing affects inflation dynamics. Furthermore, our study delves into various mechanisms and sources of heterogeneity contributing to increased inflation volatility, such as the sectors influenced by carbon pricing initiatives (e.g., the electricity sector),

⁵A carbon tax is a tax on the supply of fuel in proportion to its carbon content.

the significance of high-polluting sectors in the overall productive structure, and the role of free emission allowances in ETS schemes.

We find evidence of ETS schemes inducing larger volatility in consumer price inflation, while no significant impact is found in the case of carbon taxes. This effect feeds through the energy component, which tends to be more closely related to carbon pricing, while it affects core inflation with some delay, suggesting lagged indirect effects of ETSs on inflation. We also find that increased volatility patterns seem related to more volatile carbon prices rather than to differences in sectoral coverage among the two types of carbon pricing initiatives. Finally, our findings suggest that countries with higher weight in high-polluting sectors experience more volatile energy CPI inflation and that, under ETS schemes, a higher share of free emission allowances allocated to each country tends to reduce inflation volatility.

By shedding light on these dynamics and mechanisms, this paper adds to the growing body of literature exploring the impact of transition policies on the business cycle and on the nominal side of the economy with implications for the communication, implementation, and transmission of monetary policy (Annicchiarico et al. 2021; Carattini, Heutel, and Melkadze 2021; Diluiso et al. 2021; European Central Bank 2021; Network for Greening the Financial System 2021; Ferrari and Nispi Landi 2023; Nakov and Thomas 2023). Specifically, if carbon pricing only affects transiently the level of inflation, sharing the features of a negative supply-side shock, credible monetary policy authorities will generally “look through” this effect, as it shares the features of a temporary negative supply-side shock (Batten, Sowerbutts, and Tanaka 2020). But if carbon pricing schemes are also found to affect inflation volatility, those relative price shocks could further complicate the conduct of monetary policy, as the central bank’s signal-extraction problem becomes more challenging. Moreover, if indeed, as suggested by previous research, inflation volatility may affect the level of inflation (Cukierman and Meltzer 1986) and increase the pass-through of supply shocks to consumer prices (Arndt and Enders 2024), central banks may have to recalibrate their policies to consider those additional (upward) pressures. Finally, our results contribute to the ongoing debate on the choice of the optimal carbon pricing instrument (Parry, Black, and Zhunussova 2022; Nakov and Thomas 2023) by shedding light

on novel side effects associated with ETS initiatives compared to carbon taxes.

The rest of the paper is structured as follows. Section 2 reviews the literature. Section 3 explains the constructed database, including its descriptive statistics. Section 4 deals with the empirical model and the econometric approach considered. The results are discussed in Section 5 and, finally, Section 6 concludes.

2. Literature Review

Despite the increasing presence of carbon pricing schemes in the last decades, the literature on their economic impact has remained relatively scarce, and mostly centered on *ex ante* (model-based) evaluations. That is the case of an expanding body of research on the impact of transition policies on the business cycle, with significant contributions from scholars such as Annicchiarico et al. (2021), Carattini, Heutel, and Melkadze (2021), or Diluiso et al. (2021). Regarding the *ex post* economic impact of carbon pricing, a growing body of research has recently addressed the study of their effects on economic growth (Metcalf and Stock 2023), industrial and business performance, competitiveness and innovation (Martin et al. 2014; Martin, Muûls, and Wagner 2016; Venmans, Ellis, and Nachtigall 2020; Shapiro and Metcalf 2021), employment (Martin et al. 2014; Yamazaki 2017), international trade (Mundaca, Strand, and Young 2021), and inequality of income distribution (Elkins and Baker 2001; Känzig 2023).

With respect to inflation, in recent years the field has gained some momentum. In 2014, McKibbin, Morris, and Wilcoxon (2014) derived model-based evaluations of the introduction of a carbon tax in the U.S., pointing to sizable positive and temporary effects on inflation. Some recent work has also looked at inflationary effects of carbon pricing using New Keynesian frameworks (see, for example, Annicchiarico and Di Dio 2017; Economides and Xepapadeas 2018; Del Negro, di Giovanni, and Dogra 2023), generally finding that transition policies may induce temporary effects on inflation. The effects of carbon taxes on inflation volatility have also been addressed by means of environmental DSGE models, suggesting that disorderly scenarios in which climate policies are delayed may generate an increase in inflation volatility (Diluiso et al. 2021).

Regarding ex post analysis of the impacts on inflation, a recent study by Konradt and Weder di Mauro (2021) concludes that carbon taxes in Europe and Canada have not been inflationary, and they may even have been deflationary. They stress that the observed inflationary pressures seem restricted to the headline component and do not affect core inflation, which can be attributed to the fact that carbon taxes alter relative prices—by increasing the price of energy—but do not necessarily lead to increases in the price of the broader basket of goods and services. Focusing on the euro area, McKibbin, Konradt, and Weder di Mauro (2021) further explore this relationship, finding a positive effect of carbon taxes on inflation, particularly in the first years following the introduction of the carbon tax. The effect is nevertheless found to fade in the medium to long term. Moreover, the impact on core inflation tends to be negative, also pointing to carbon taxes affecting inflation through the change in relative prices instead of affecting overall price levels. By exploiting high-frequency data and some institutional features of the European carbon market, Känzig (2023) also derives a positive impact of carbon pricing on energy and consumer prices. Indeed, carbon policy shocks are found to account for around one-third of the variations in energy prices in the context of the EU ETS. In a similar fashion, Moessner (2022) analyzes the effects of carbon pricing on the level of headline, food, energy, and core inflation for OECD countries, focusing not only on carbon taxes but also on the prices of ETSs. She finds that an increase in the price of ETSs is indeed inflationary but its effects are circumscribed to energy CPI inflation and fade away in two years. Surprisingly, an increase in carbon taxes is found to push up food CPI inflation only, and it shows no significant effects on the rest of components.

To date, this body of research has focused on exploring the ex post impact of carbon pricing on the level of inflation. By empirically assessing the effects of carbon pricing on inflation volatility, our paper speaks to several strands of the literature. First, we aim to contribute to the long-lived debate on prices versus quantities (Weitzman 1974) and the ongoing discussion on the choice of the optimal carbon pricing instrument (see Parry, Black, and Zhunussova 2022 for a review). Carbon pricing systems can take the form of both carbon taxes and cap-and-trade systems, with presumed differing impacts on inflation volatility. As outlined in the introduction,

carbon taxes tend to deliver relatively stable carbon prices, while ETSs have been found to lead to substantial carbon price volatility (Goulder and Schein 2013; Metcalf 2021). Higher volatility of carbon prices could be transmitted to producer and consumer prices—especially of those goods and services related to those sectors under the coverage of carbon pricing—and thus increase inflation volatility.⁶ Second, we contribute to the existing literature on the impact of transition policies on the business cycle (Annicchiarico et al. 2021; Carattini, Heutel, and Melkadze 2021; Diluiso et al. 2021). Our work adds to this discussion by identifying another potential source of macroeconomic volatility. Finally, our paper is connected to the emerging body of research on how climate change and mitigation policies might influence central banks' operations and the ability of monetary policy to pursue price stability. Noteworthy contributions in this area include those by Carattini, Heutel, and Melkadze (2021); Diluiso et al. (2021); Ferrari and Nispi Landi (2023); Kaldorf and Giovanardi (2023); and Nakov and Thomas (2023).

3. Data

Our sample starts in 2000, after the introduction of the euro⁷ and just before the accession of China to the World Trade Organization,⁸ two major events affecting the inflation process in advanced economies. We focus our analysis on OECD countries, as inflation in advanced economies has behaved differently than in emerging market economies (Blanchard, Cerutti, and Summers 2015; Ha, Kose, and Ohnsorge 2019). This results in an unbalanced panel of 36 countries for the period 2000–20.

⁶Nevertheless, it should also be noted that the introduction of carbon price initiatives could also reduce the weight of fossil fuels in the CPI consumption basket (Andersson 2019), which has historically been a source of volatility of inflation.

⁷That is because a substantial proportion of OECD countries are part of the euro area. The creation of the European Monetary Union in 1999 led to structural changes in monetary policy and inflation patterns in those countries (Bowdler and Malik 2005; Balatti 2020).

⁸China's accession to World Trade Organization in 2001 is often seen a milestone in globalization. The globalization of inflation hypothesis argues that the factors influencing inflation dynamics are becoming increasingly global (Forbes 2019).

Our dependent variable for inflation volatility is constructed from CPI data for headline, core, and energy inflation, which are retrieved from the OECD Data Portal with monthly frequency. Inflation volatility for these three indices is computed as the standard deviation of year-on-year inflation for each country in windows of 24 months.⁹ Also, as a robustness test, in order to avoid unduly overweighting extreme data points, following a common approach in the literature, inflation volatility is computed as the log of one plus the standard deviation of inflation (Bowdler and Malik 2017).¹⁰ There is also evidence that carbon pricing leads to differing impacts according to the component of inflation considered, i.e., headline, energy, or core inflation (McKibbin, Konradt, and Weder di Mauro 2021; Moessner 2022; Känzig 2023; Konradt and Weder di Mauro 2023), which could translate to differential impacts on the volatility of that inflation. Therefore, we compute these measures for headline, energy, and core inflation separately.

Information on carbon pricing comes from the Carbon Pricing Dashboard, a comprehensive database with global geographical coverage elaborated by the World Bank, that includes data on either national, subnational, and supranational carbon pricing initiatives.¹¹ By processing the information contained in the database, we build some variables of interest for this work. First, the economic impact of carbon pricing has proved to depend on the level of coverage of the initiatives, on the basis that the pass-through to consumer prices is expected to be higher if the initiative covers a greater proportion of emissions within its jurisdiction (Metcalf and Stock 2023). The effect of carbon pricing on inflation should also depend on the price of the initiative (Konradt and Weder di Mauro 2021, 2023; Metcalf and Stock 2023). Therefore, following Konradt and Weder di Mauro (2023) and Metcalf and Stock (2023), we compute the real effective

⁹The choice of the frequency was motivated by the need to compute inflation volatility relying on a sufficient number of observations. However, it should be noted that it led to the exclusion of Australia and New Zealand, for which only quarterly data are available. Figure A.1 in Appendix A depicts the evolution of this measure of inflation volatility.

¹⁰This allows to downweight very large readings of inflation that may occur during hyperinflation episodes as well as inflation records close to zero.

¹¹See <https://carbonpricingdashboard.worldbank.org>.

carbon price in each jurisdiction for each year (*Price*), weighting the carbon price of each initiative within the same country by the share of emissions covered by each initiative.¹²

Additionally, the effects of carbon pricing systems on inflation may depend on some other factors, which we control for. Carbon pricing initiatives may coexist with more stringent environmental norms, which may impose some implicit abatement costs alongside the explicit price for carbon. The role of those other climate change mitigation policies is captured from the OECD's Environmental Stringency Index,¹³ which aims to measure environmental policy stringency internationally over a relatively long time horizon. Finally, to consider the possibility of a structural change after the surge in environmental policies following the Paris Agreement in 2015, we create an additional dummy variable (*Post-2015*). Additional potential mechanisms through which carbon pricing may affect inflation are explored in Section 5.3.

Regarding control variables, the literature has identified a number of domestic and global factors that affect inflation volatility. First, with respect to domestic factors, the size of the country and its level of development are usually considered as drivers of inflation volatility (Lane 1997; Neely and Rapach 2011; Parker 2018). In addition, the domestic level of CPI inflation has also been identified as a factor affecting its volatility (Bleaney and Fielding 2002). Second, some domestic factors related to the sensitivity of the economy to external conditions are also commonly associated with inflation volatility. For example, trade, its role as commodity exporter, or the *de iure* and *de facto* capital account liberalization of each economy are all related to the reaction of domestic prices to swings in external economic conditions (Romer 1993; Lane 1997; Aghion, Bacchetta, and Banerjee 2004; Bowdler and Malik 2017). Third, a strand of the literature has found that inflation dynamics are increasingly explained by global factors, and particularly international commodity prices, rather than by domestic conditions (Fernández,

¹²Emissions not covered by any initiative are supposed to have a price of zero; therefore, the variable *Price* could be understood as a sort of real effective price of CO₂ in each country. Carbon prices are deflated by the national CPI and expressed in 2018 dollar terms.

¹³The correlation between having carbon pricing initiatives and environmental stringency is indeed not as high as could be expected, yielding a value of 0.56.

Schmitt-Grohé, and Uribe 2017; Parker 2018; Kamber and Wong 2020).

In order to account for those factors, we gathered information from different sources. Population, GDP per capita, and trade openness are directly drawn from the World Bank's World Development Indicators. Our proxy for capital flows and financial openness is built as the sum of the absolute value of inflows and outflows as a percentage of each country's GDP based on the International Monetary Fund's International Financial Statistics data. Using data on merchandise exports from the World Bank, a country is considered a primary exporter if exports of food, fuels, ores, and metals constitute more than 20 percent of total merchandise exports. Following Fernández, Schmitt-Grohé, and Uribe (2017) and Kamber and Wong (2020), real commodity prices are accounted for through the growth rate in Agricultural, Metal and Minerals, and Energy Commodity Price Indices, deflated by the U.S. CPI, obtained from the World Bank Pink Sheet on commodity prices.

Table 1 shows the main descriptive statistics of our annual panel for all relevant variables.

With respect to carbon pricing, Figure 1A shows that the number of initiatives implemented in OECD countries has grown sharply in the last decades, going from only 7 initiatives in 2000 to 45 initiatives in 2020. Coverage has also risen steadily. While in 2000 existing OECD initiatives covered together only 0.45 percent of global emissions, by 2020 the coverage of global emissions from OECD carbon pricing initiatives reached 11 percent. As shown in Figure 1B, on a global scale, the evolution of the number of initiatives and coverage followed a similar pattern during the period, with 58 initiatives already implemented worldwide by 2020 and a global coverage of 15 percent of emissions. It is also worth noting that OECD countries involve a substantial proportion of all initiatives implemented globally; by considering OECD countries, we are accounting for 78 percent of initiatives and 73 percent of total emissions covered by those initiatives.

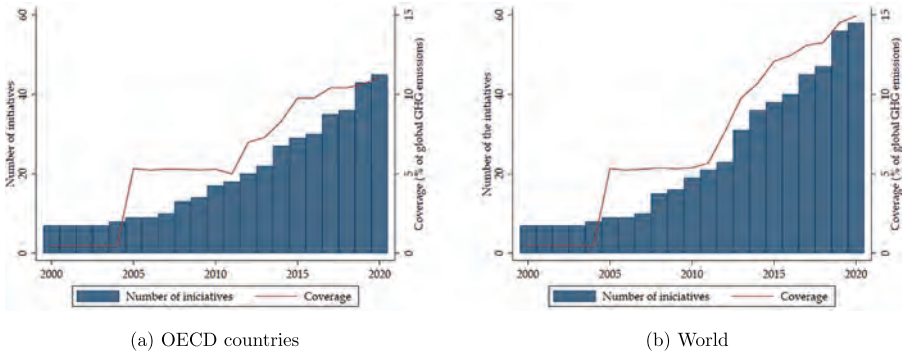
Figure 2 further depicts a timeline of the introduction of carbon pricing in OECD countries for the period 2000–20. As can be observed, our sample shows substantial variability both in geographical terms and across time, with initiatives of different types—ETSs and carbon taxes—implemented in a number of countries and at

Table 1. Descriptive Statistics

		N	Mean	SD	Min.	Max.
Inflation Volatility	Headline Infl. Standard Deviation (Percentage Points)	756	0.981	1.173	0.159	15.40
	Energy Infl. Standard Deviation (Percentage Points)	756	5.331	3.765	0.709	38.70
	Core Infl. Standard Deviation (Percentage Points)	756	0.741	1.023	0.104	12.91
Carbon Pricing	Real Price (2018 US\$/tCO ₂ e)	756	4.641	10.15	0	89.07
	Real Price (ETS) (2018 US\$/tCO ₂ e)	756	4.086	5.366	0	20.33
	Real Price (Carbon Tax) (2018 US\$/tCO ₂ e)	756	4.067	12.88	0	103.3
	Environmental Stringency, 0 (Not Stringent) to 6 (Highest Degree)	672	2.516	1.009	0	4.890
	Post-2015, Dummy	756	0.238	0.426	0	1
Control Variables	GDP per capita (in Thousands)	756	39.94	18.24	9.047	113.8
	Population (in Millions)	756	35.07	56.45	0.281	329.5
	Trade Openness (% of GDP)	754	96.06	57.54	19.80	408.4
	Primary Exporter (Dummy)	756	0.431	0.496	0	1
	Capital Flows (% of GDP)	718	88.04	365.9	2.670	3,607
	Commodity Prices: Energy Inflation (%)	756	6.596	25.49	-45.13	58.32
	Commodity Prices: Agriculture Inflation (%)	756	1.190	9.741	-13.82	22.98
	Commodity Prices: Metals and Minerals Inf. (%)	756	4.962	19.73	-22.85	48.47
	CPI Inflation (%)	756	2.699	3.910	-4.990	68.53

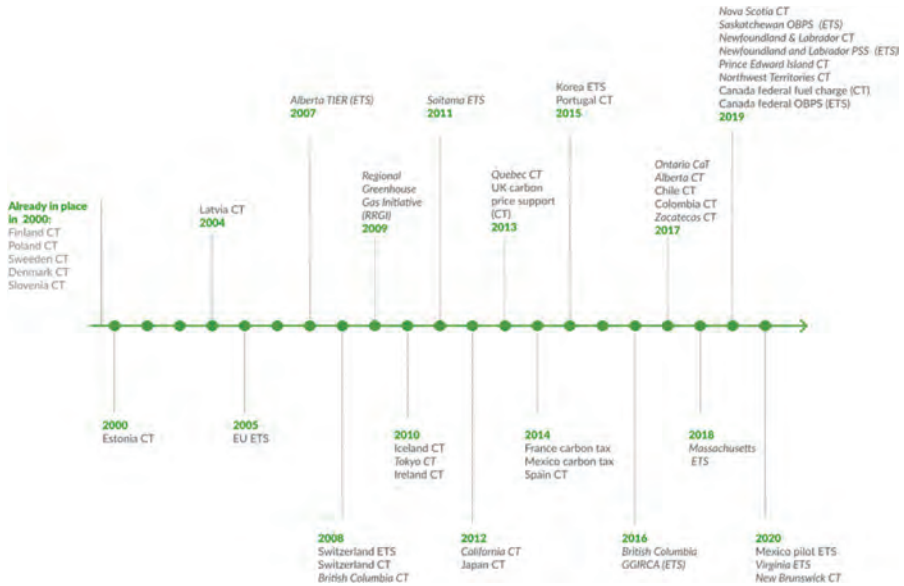
Note: N: number of observations; SD: standard deviation.

Figure 1. Number of Carbon Pricing Initiatives and Coverage



Source: Authors’ own elaboration based on World Bank Carbon Pricing Dashboard.

Figure 2. Timeline for Initiatives of Carbon Pricing in OECD Countries



Source: Authors’ own elaboration based on World Bank Carbon Pricing Dashboard.

different points in time. Those initiatives are also heterogeneous in terms of their coverage and price for carbon.

4. Empirical Strategy

In our analysis, we employ a widely utilized empirical framework from the literature to assess the economic impacts of carbon taxes, which has been employed by previous studies (Konradt and Weder di Mauro 2021; Metcalf and Stock 2023). Specifically, we use panel local projections (Jordà 2005) as our central specification.¹⁴

To estimate the dynamic response of CPI inflation volatility at horizon h to a carbon price shock, we adapt the specifications introduced by Konradt and Weder di Mauro (2023) and Metcalf and Stock (2023):

$$\begin{aligned} \sigma(\pi)_{i,t+h} = & \Theta^h \text{Carbon price}_{i,t} + \beta^h(L) \text{Carbon price}_{i,t-1} \\ & + \delta^h(L)\sigma(\pi)_{i,t-1} + \gamma^h X_{i,t-1} + \mu_i^h + \epsilon_{i,t+h}, \end{aligned} \quad (1)$$

where $\sigma(\pi)_{i,t}$ is the standard deviation of monthly year-on-year inflation for a country i and for the year t , computed in windows of 24 months. We consider the standard deviation of headline, energy, and core inflation. $\text{Carbon price}_{i,t}$ is the real effective carbon price, for which we distinguish among those associated with any carbon pricing initiative, to ETSs and to carbon taxes. Θ^h can be interpreted as the effect of an unexpected change in the carbon price in year t on inflation volatility in h years, under the identifying assumption that a component of the effective carbon price is not predictable by past macroeconomic controls. $X_{i,t-1}$ includes all the relevant control variables described in Section 3.^{15,16} μ_i^h are country fixed

¹⁴Please note that an alternative approach, utilizing panel data models, was followed in a previous working paper version of this study. For details on this alternative approach, please refer to Appendix B.

¹⁵In our setting, some of the controls may be contemporaneously affected by the carbon tax. For instance, commodity prices are both an outcome and a channel through which carbon prices affect inflation, and contemporaneous GDP could be affected by carbon taxes. Therefore, we include lagged control variables. However, in Section 5.4 we show that the results are robust to the inclusion of contemporaneous control variables.

¹⁶We do not include time fixed effects because, as noted by Känzig and Konradt (2023), when examining supranational carbon pricing initiatives such as the

effects and $\epsilon_{i,t+h}$ represents the error term. $\beta^h(L)$ and $\delta^h(L)$ represent lag polynomials, reflecting the use of the four latest lags to control for persistence of the carbon pricing and inflation volatility, respectively.


Our identification approach thus hinges on several strategies. First, as in Metcalf and Stock (2020) and Konradt and Weder di Mauro (2021, 2023), we rely on time-series identification. That is, by using an extensive set of controls, as well as past values of inflation volatility and carbon pricing, we can identify the impact of changes in carbon prices that are not predicted by past economic conditions, or unobserved historical factors that could simultaneously affect inflation volatility and carbon prices. The remaining variation can thus be assumed to be unexpected. Second, although we should not expect endogeneity issues within this framework, in Section 5.5 we test the robustness of our results to an instrumental variables (IV) approach using panel local projections IV (LP-IV) along the lines of Jordà, Schularick, and Taylor (2015), using a series of carbon policy surprises developed by Känzig (2023) as our primary instrument for carbon prices.

The impulse responses are constructed based on the fixed effects estimates of Θ^h coefficients at each horizon. As in Konradt and Weder di Mauro (2021, 2023) and Metcalf and Stock (2023), the shock consists in a one-time permanent increase in the effective carbon price by USD 40 applied to 30 percent of each economy's GHG emissions, equivalent to an increase of the effective carbon price of USD 12.¹⁷ The reported 95 percent confidence bands are based on

European carbon market, it becomes impractical to include time fixed effects. This is because the policy variation of interest primarily occurs at the supranational level, and including time fixed effects would absorb the majority of this relevant variation. Moreover, time fixed effects are collinear with some of the control variables, namely, the ones related to global factors: energy commodity prices inflation, agriculture commodity prices inflation, and metals and minerals commodity prices inflation. Nevertheless, a robustness test to the inclusion of time fixed effects is included in Section 5.4.

¹⁷As we mentioned, the effect of carbon pricing systems on inflation volatility may depend on the amount of GHG emissions covered by carbon pricing initiatives in each country and the carbon price paid in this jurisdiction.

Figure 3. Impulse Responses of Headline Inflation Volatility to Shocks in Carbon Prices



Note: The solid blue lines describe the dynamic impulse responses to an increase in the effective carbon price by USD 40 with 30 percent emission coverage. The gray area delimits 95 percent confidence bands. All responses are estimated by local projections including country fixed effects and controlling for GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices: energy inflation, agriculture inflation, and metals and minerals inflation; and CPI inflation. Standard errors are heteroskedasticity robust and clustered by country.

the respective estimated standard errors clustered by country and are heteroskedasticity robust.

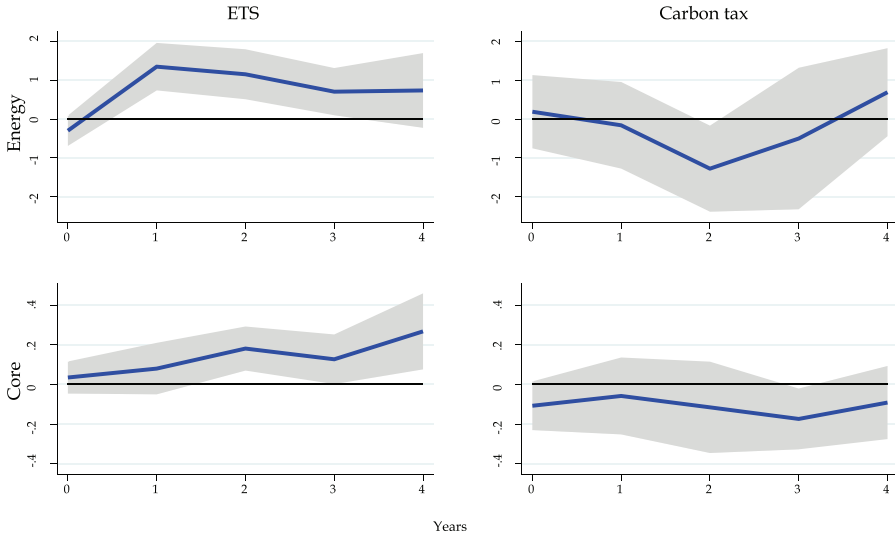
5. The Impact of Carbon Pricing on Inflation Volatility

5.1 Baseline Results

Our baseline results presented in Figure 3 depict, respectively, the effects of a rise in effective carbon prices on headline inflation volatility without further differentiating by types (left-hand-side chart), and considering the ETS and carbon tax types separately (center and right-hand-side charts).

Overall, a preliminary analysis suggests that the presence of a carbon pricing system is associated with increased inflation volatility, although the pattern is not entirely clear. However, when differentiating between ETS schemes and carbon taxes, distinct patterns emerge. An increase in ETS effective prices is linked to a rise in the volatility of headline inflation, whereas carbon tax rates do not have a significant impact and are only found to marginally reduce volatility in one of the years.

Figure 4. Impulse Responses of Energy and Core Inflation Volatility to Shocks in ETS Prices and Carbon Tax Rates



Note: The solid blue lines describe the dynamic impulse responses to an increase in the effective carbon price by USD 40 with 30 percent emission coverage. The gray area delimits 95 percent confidence bands. All responses are estimated by local projections including country fixed effects and controlling for GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices: energy inflation, agriculture inflation, and metals and minerals inflation; and CPI inflation. Standard errors are heteroskedasticity robust and clustered by country.

5.2 Channels: Effects on the Volatility of Energy and Core Inflation

In a second step, we investigate the channels by which carbon pricing systems affect inflation volatility, analyzing the differential impact on the energy CPI component and the degree of pass-through to the volatility of core inflation. For this purpose, in Figure 4 the dependent variables are energy and core inflation volatility, respectively, and the effects of ETSs and carbon taxes are considered separately. Control variables from the baseline specification are also included.

The estimates reveal that only schemes of the ETS type push up the volatility of the energy and core components of inflation,

while overall such impact is not found in the case of carbon taxes.¹⁸ ETS schemes show a significant impact on inflation volatility for the case of the energy component, which should be expected given that energy-related goods are the ones that are targeted directly or indirectly by carbon pricing initiatives. The volatility of core inflation is found to be affected with some delay, after two years, in those jurisdictions with ETSs, which suggests lagged indirect effects of ETS schemes on core inflation volatility.

These results align with prior research indicating that the implementation of carbon pricing schemes affects inflation primarily through the energy component (Konradt and Weder di Mauro 2021; McKibbin, Konradt, and Weder di Mauro 2021; Moessner 2022). It is also in line with evidence on the indirect effects of other cost-push shocks, such as rising energy prices, which may lead to increased production costs and subsequent pass-through to consumer prices. In fact, Kilian and Zhou (2023) found that energy price shocks affect overall U.S. inflation through the energy component of the consumer basket, but also exhibit effects on core inflation. For the euro area, López, Párraga Rodríguez, and Santabárbara (2022) and Adolfsen et al. (2024) show a significant transmission of gas prices (due to their role in electricity pricing) to all inflation components, including core inflation, albeit with some lag.

5.3 Mechanisms: Sectoral Coverage, Production Structure, Volatility of ETSs, and ETS Free Allowances

In this section, we investigate potential mechanisms driving the effects of carbon pricing on inflation volatility, such as the sectoral coverage of carbon pricing initiatives, the differences in countries' production structure, the volatility of ETS prices, or the amount of ETS pollution rights allocated freely. These factors collectively

¹⁸In fact, carbon taxes not only fail to increase inflation volatility, but they are actually found to slightly reduce it for one of the years. This reduction can be attributed to two main factors. Firstly, the implementation of carbon pricing can lead to a reduction in the relative weight of fossil fuels in the CPI consumption basket, historically a significant source of inflation volatility (Andersson 2019). Secondly, an increase in the carbon tax implies a larger share of the fixed component of the energy price, potentially mitigating the volatility of energy consumer prices in response to fluctuations in energy input prices.

contribute to explaining the diverse effects observed in carbon pricing initiatives. To capture these effects, we adapt the local projection baseline specification by including an interaction term that considers each of these mentioned mechanisms.

$$\begin{aligned} \sigma(\pi)_{i,t+h} = & \Theta^h \text{Carbon price}_{i,t} + \alpha^h \text{Carbon price}_{i,t} \\ & \times \text{mechanism}_{i,t} + \beta^h(L) \text{Carbon price}_{i,t-1} \quad (2) \\ & + \delta^h(L)\sigma(\pi)_{i,t-1} + \gamma^h X_{i,t-1} + \mu_i^h + \epsilon_{i,t+h}, \end{aligned}$$


where α^h captures the differential impact of carbon prices according to the mechanism considered. As in Känzig and Konradt (2023), variables that capture these mechanisms are standardized so coefficients can be interpreted as the impact of changes corresponding to one standard deviation compared to the average country.

5.3.1 Carbon Pricing Sectoral Coverage

The finding that ETS schemes are associated with larger inflation volatility could potentially stem from two plausible explanations. Firstly, a contributing factor could be the divergent coverage of cap-and-trade systems and carbon taxes. For instance, electricity producers tend to be encompassed within ETSs, while in the case of taxes, they more frequently exclude the electricity generation sector. These distinct coverage scopes may influence the impact on inflation volatility. Second, ETS prices, being market driven, inherently display higher levels of volatility compared to carbon taxes, which are set at less frequent intervals. This increased volatility in cap-and-trade prices could potentially translate into higher overall inflation volatility. By further examining these factors, we can gain deeper insights into the mechanisms driving the differential effects of cap-and-trade schemes and carbon taxes on inflation volatility.

Starting with the divergent coverage of ETSs and carbon taxes, one potential approach to explore this hypothesis is to examine whether differences between ETS prices and carbon taxes persist when analyzing regimes with similar coverage, particularly those that include the electricity generation sector. By narrowing the focus to these specific regimes, we can gain insights into whether the observed differences in price dynamics between cap-and-trade and

Figure 5. Effects of Carbon Pricing Depending on Electricity-Sector Coverage (only national carbon taxes)



Note: Impulse responses to an increase in effective carbon taxes of USD 40 with 30 percent emission coverage interacted to a variable indicating whether electricity sector is covered by the system. The gray area delimits 95 percent confidence bands. All responses are estimated by local projections including country fixed effects and controlling for GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices; energy inflation, agriculture inflation, and metals and minerals inflation; and CPI inflation. Standard errors are heteroskedasticity robust and clustered by country.

carbon taxes are primarily driven by variations in sectoral coverage. Unfortunately, all national and supranational ETSs within our sample cover the electricity sector, with the notable exception of Canada’s federal ETS, which was only introduced in 2019. Therefore, we cannot compare across ETSs including and excluding the electricity sector. Nevertheless, regarding carbon taxes, our analysis incorporates a significant number of carbon taxes both with and without the inclusion of the electricity sector (see Table A.4), which allows us to test this hypothesis.

As shown in Figure 5, no significant differences in volatility are encountered for carbon taxes that include and exclude the electricity sector. This supports the hypothesis that the coverage of the electricity sector may not be the primary driver of the observed disparities in inflation volatility between ETSs and carbon taxes.

5.3.2 Volatility in ETS Prices

In a subsequent step, we extend our analysis to investigate the implications of their differing level of carbon price volatility in the cap-and-trade regimes, as discussed earlier. This approach allows

us to gain a deeper understanding of the role of differences in carbon price volatility as a mechanism driving the divergent impact of ETSs and carbon taxes on inflation volatility. To accomplish this, we obtained the price for the major national and supranational cap-and-trade regimes included in our sample and calculated their respective volatility.¹⁹

The results indicate that headline inflation, energy inflation, and core inflation exhibit higher levels of volatility in response to increased volatility in ETS prices (see Figure 6).²⁰ That is, ETS schemes tend to have a more significant impact on inflation volatility across components, as their prices are indeed more volatile. These findings therefore provide suggestive evidence that carbon price volatility may be a key mechanism driving the observed effects, and that the inherent volatility within the cap-and-trade framework plays a crucial role in shaping the relationship between carbon pricing and inflation.


5.3.3 Specialization in High-Polluting Sectors

The impact of carbon pricing can also vary depending on a country's productive structure. A given shock is expected to have a more intense effect on inflation depending on the country specialization in high-polluting industries and sectors, which we measure by the share of value-added in each economy generated by the top 20 percent most polluting sectors worldwide using OECD data.

As shown in Figure 7, both ETS and carbon tax shocks indeed appear to have a comparatively higher impact when a larger share of value-added can be attributed to high-polluting sectors, with the effect primarily observed in the energy CPI volatility.

¹⁹Those include the EU ETS, as well as the Switzerland, Mexico, and Korea ETSs. We compute the coefficient of variation to build a measure of volatility comparable across ETS regimes. The ETS prices used in the analysis are sourced from the International Carbon Action Partnership.

²⁰However, it is important to note that an exception arises in the fourth year. This deviation may be attributed to enhanced incentives within more volatile regimes to replace fossil fuels as a means to mitigate the higher impact of carbon pricing on such regimes, which may reduce the effects on inflation volatility over the medium term.

Figure 6. Effects of ETS Price Volatility

Note: Impulse responses to an increase in effective ETS prices of USD 40 with 30 percent emission coverage interacted to the volatility of carbon price within ETS regimes. The gray area delimits 95 percent confidence bands. All responses are estimated by local projections including country fixed effects and controlling for GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices; energy inflation, agriculture inflation, and metals and minerals inflation; and CPI inflation. Standard errors are heteroskedasticity robust and clustered by country. The ETS prices used in the analysis are sourced from the International Carbon Action Partnership. Volatility is measured using the coefficient of variation due to the substantial dispersion of prices within the included regimes.

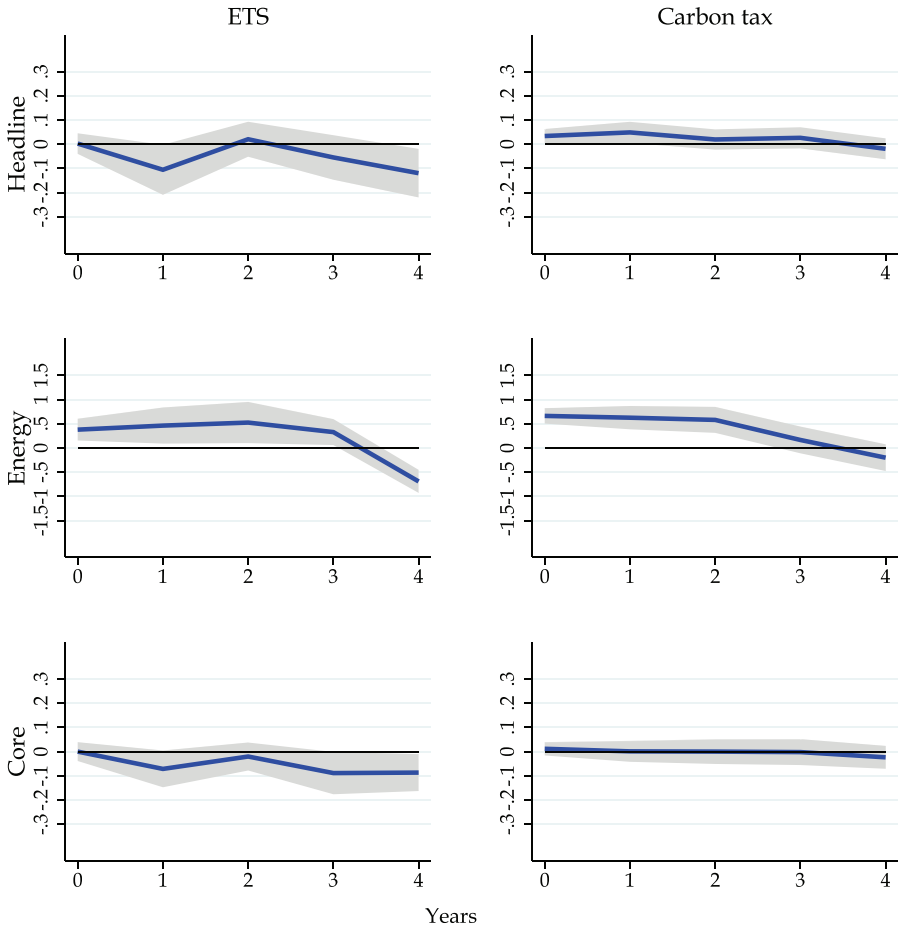
5.3.4 ETS Free Allowances

Finally, another potential avenue to explore the mechanisms behind the effects of carbon pricing on inflation volatility is the allocation of free allowances within cap-and-trade schemes, which has shown substantial variation across EU countries and time. We focus our analysis on the EU ETS, for which we have available data.²¹

The allocation of free allowances has been seen as a strategic policy choice aimed at addressing concerns such as potential loss of competitiveness and carbon leakage. By providing free allowances to certain industries or sectors, policymakers aimed at alleviating the immediate burden of carbon pricing on these entities. This allocation mechanism effectively reduces the costs they would otherwise incur from purchasing allowances, thereby lowering the effective carbon price they bear.


²¹Data on free allowances are taken from the EU Transaction Log, being available from 2005, the year of the establishment of the EU ETS.

Figure 7. Effects of Carbon Pricing by Value-Added Generated by Polluting Sectors



Note: Impulse responses to an increase in effective carbon prices of USD 40 with 30 percent emission coverage interacted to the value-added generated by polluting sectors. The gray area delimits 95 percent confidence bands. All responses are estimated by local projections including country fixed effects and controlling for GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices: energy inflation, agriculture inflation, and metals and minerals inflation; and CPI inflation. Standard errors are heteroskedasticity robust and clustered by country.

Figure 8. Effects of Carbon Pricing by Share of Free Allowances in the EU ETS Market



Note: Impulse responses to an increase in effective EU ETS prices of USD 40 with 30 percent emission coverage interacted to the share of free allowances. The gray area delimits 95 percent confidence bands. Standard errors are heteroskedasticity robust and clustered by country. Data on free allowances sourced from EU Transaction Log.

Interestingly, our findings show that a higher share of free allowances is associated with lower effects of ETS prices on inflation volatility, with the exception of the energy component, where no significant effect was found (see Figure 8). This supports the hypothesis that higher allocation of free allowances may help prevent price volatility. This outcome is not surprising when considering the historical functioning of the EU ETS. At the start of the current trading period, manufacturing industries received 80 percent of their allowances for free, while the electricity sector—the one directly related to the energy component—has not received any free allowances since 2013.²²

5.4 Robustness Tests

Finally, in this section we present further evidence of our results being robust to the use of different specifications, subsamples, and additional controls.

²²See https://climate.ec.europa.eu/eu-action/eu-emissions-trading-system-eu-ets/free-allocation_en#documentation for more detailed information on the functioning of the allocation of free allowances across industries and sectors.

First, as described in Section 3, we could expect that the existence of a structural break after the renewed surge in environmental policies following the Paris Agreement in 2015 could be driving the results. Tables A.1 and A.2 in Appendix A show results restricting the sample to the years before the Paris Agreement, finding no evidence of those affecting the results. Also, our findings prove robust to the inclusion of the OECD index of environmental stringency, in an attempt to control for the fact that our variables of carbon pricing could be capturing the effects of changes in each country's environmental stringency not necessarily related (see again Tables A.1 and A.2 in Appendix A). Second, we add a specification in which time dummies are included, at the expense of removing those variables collinear with them, showing that main results hold under this alternative modeling choice (Tables A.1 and A.2 in Appendix A). Third, we test our specification against the use of the log of one plus the standard deviation of inflation, as described in Section 3. As shown in Figure A.2 in Appendix A, our main results are robust. Fourth, we show that the results are robust to the inclusion of contemporaneous (instead of lagged) control variables (see Figure A.3).

5.5 *IV Approach: Local Projections IV*

While carbon pricing should not be expected subject to endogeneity in our baseline empirical framework, this section examines the robustness of our findings to an IV approach.

For that purpose, we employ panel local projections IV (LP-IV) along the lines of Jordà, Schularick, and Taylor (2015). Dynamic causal effects are estimated from the same set of equations as in (2), where *Carbon price*_{*i,t*} is the possible endogenous regressor and *z*_{*t*} acts as an instrument for *Carbon price*_{*i,t*}. As in the baseline specification, we also include all relevant lagged controls, four lags of carbon pricing and inflation volatility, and country fixed effects. Errors are heteroskedasticity robust and clustered by country.

5.5.1 *Instrumental Variable: Carbon Policy Surprises*

As our primary instrument, we rely on a set of EU ETS carbon policy surprises developed by Känzig (2023). However, this approach requires us to reduce our sample to the countries covered by the EU

ETS. This study builds upon the existing body of literature on high-frequency identification techniques, which identify policy surprises by analyzing the movements of asset prices following significant policy events. While commonly applied in the context of monetary policy (Kuttner 2001; Gürkaynak, Sack, and Swanson 2005; Gertler and Karadi 2015; Nakamura and Steinsson 2018) or global oil markets (Känzig 2021), Känzig (2023) adapts this approach to the domain of carbon pricing by applying it specifically to the EU ETS market.

To construct a series of carbon policy surprises, Känzig (2023) compiles a comprehensive list of regulatory events related to changes in the supply of emission allowances within the EU ETS.²³ This list encompasses various events related to, among others, decisions made by the European Commission, votes conducted by the European Parliament, and judgments rendered by European courts. In total, 126 such events are identified, covering the period from 2005 to 2019. Carbon policy surprises are then computed as the change in the EU emission allowances' futures prices on the day of the regulatory event compared to the last trading day before the event, in percentage terms or measured relative to the prevailing wholesale electricity price on the day before. The monthly series is constructed by aggregating the daily surprises, adding them up over each month. In cases where no regulatory events take place during a particular month, the series is assigned a value of 0. Yearly series are constructed by adding up monthly series.

$$CPSurprise_{t,d} = \frac{F_{carbon,t,d} - F_{carbon,t,d-1}}{F_{carbon,t,d-1}} \cdot 100 \quad (3)$$

We take the monthly surprises in percentage of ETS prices and aggregate them yearly to match our specification.

As highlighted by Stock and Watson (2018), variables based on high-frequency surprises such as the ones we use in this section are susceptible to measurement error, since they may not capture all relevant events. While they may still serve as valid external instruments, employing them directly can yield biased estimates of the

²³That allows us to address concerns that the events are capturing other information pertaining economic activity or broader factors affecting the demand of emission allowances.

Table 2. First-Stage Regressions of ETS Effective Carbon Price on the Instrument

	(1) Headline	(2) Energy	(3) Core
Carbon Policy Surprises	0.27*** (0.01)	0.24*** (0.01)	0.26*** (0.01)
F-statistic	2576.45***	2331***	3033***
Observations	586	585	584

Note: *p < 0.10, **p < 0.05, ***p < 0.01. First-stage regressions include country fixed effects and controls for GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices: energy inflation, agriculture inflation, and metals and minerals inflation; and CPI Inflation—as well as four lags of inflation volatility (headline, energy, or core depending on the equation) and effective ETS carbon prices. Country-cluster robust standard errors are reported in parentheses.

dynamic causal effects. Thus, we use them as instruments using LP-IV rather than as a direct measure of the specific shock of interest. Unfortunately, carbon policy surprises are only provided for the EU ETS market, so we will only provide LP-IV results pertaining to shocks in ETS prices.


5.5.2 Results of LP-IV

Table 2 presents the results of first-stage IV regressions. As expected, carbon policy surprises are positive and significantly correlated with carbon prices. High values for the F-statistic provide additional evidence of the significance of the instrument in the first-stage regression at the 1 percent level. The instrument's relevance is further confirmed by the Kleibergen and Paap (2006) tests of weak instruments, which are reported in Table A.3 in Appendix A for each variable and horizon.

Exogeneity conditions²⁴ cannot directly be tested in our framework. Nevertheless, the high-frequency identification strategy should be able to isolate the exogenous component of ETS carbon price shocks. Also, as discussed in Känzig (2023), the series of surprises

²⁴In an LP-IV framework, both contemporaneous and lead-lag exogeneity must hold for validity.

Figure 9. LP-IV Impulse Responses of Headline, Energy, and Core Inflation Volatility to Shocks in ETS Prices



Note: Impulse responses to an increase in effective ETS prices of USD 40 with 30 percent emission coverage. The gray area delimits 95 percent confidence bands. All responses are estimated by local projections IV using ETS carbon pricing surprises from Känzig (2023) as instrumental variable, and including country fixed effects and controlling for GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices; energy inflation, agriculture inflation, and metals and minerals inflation; and CPI inflation. Standard errors are heteroskedasticity robust and clustered by country.

used as instruments are not autocorrelated or forecastable by past macroeconomic and financial variables or other structural shock measures such as oil demand, uncertainty, financial, fiscal, and monetary policy shocks (see Appendix B.1 in Känzig 2023). Moreover, in our setting lagged macroeconomic variables should be able to control for main past macroeconomic shocks affecting inflation volatility (Stock and Watson 2018).

Figure 9 reports LP-IV impulse responses of headline, energy, and core inflation volatilities to shocks in ETS prices. As shown, our main results seem robust to the use of an IV approach. The effects of ETS effective carbon prices on energy and headline inflation volatilities are found to be qualitatively similar and of comparable magnitude, although with somewhat smoother impulse responses. Notably, the indirect effects on core inflation volatility seem to manifest with shorter delays, becoming evident as early as the first year.

6. Conclusions

In the following years, climate change and related mitigation policies are expected to affect potential growth, business cycle, and inflation

dynamics. This may reshape the formulation of economic policy, and has prompted the urge from policymakers, including central banks, to understand their economic consequences. Carbon pricing initiatives, aimed at increasing the relative prices of GHG-intensive goods and services, could not only push up inflation, but also affect inflation volatility. This paper empirically assesses the effects of carbon pricing on inflation volatility, disaggregating by the two main types of initiatives: ETSs and carbon taxes.

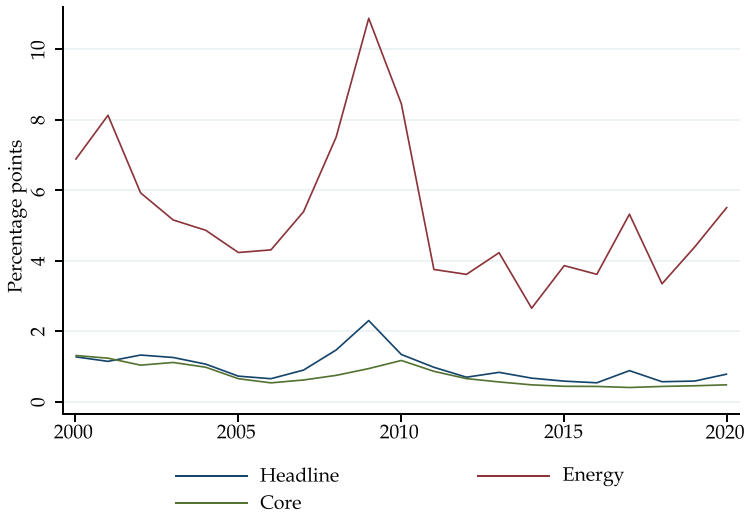
Our main findings could be summarized as follows. First, this work finds evidence of ETS schemes inducing greater volatility in consumer price inflation, while no significant impact is found in the case of carbon taxes. Second, this effect feeds through the energy component, which tends to be more closely related to carbon pricing initiatives, while it affects core inflation with some delay, suggesting lagged indirect effects of ETSs on inflation. Altogether, these results seem consistent with the fact that carbon pricing initiatives of the ETS type tend to lead to more volatile carbon prices (and carbon price inflation) than the more stable carbon taxes, and that volatility is fed into consumer prices, particularly through the energy component. It also confirms the results of previous studies that carbon pricing initiatives, by changing relative prices and increasing the price of fossil fuels, lead to a direct effect on inflation by affecting the energy component, and provide new evidence of indirect effects in the rest of goods and services included in core CPI.

Moreover, we explore several mechanisms that could drive the differences between ETSs and carbon taxes. Our findings indicate that the increased volatility observed in ETS schemes is primarily due to higher fluctuations in carbon prices, rather than to differences in the sectoral coverage compared to carbon taxes. Furthermore, we identify several factors that contribute to the heterogeneous impacts of carbon pricing initiatives across jurisdictions. For example, countries with higher weight in high-polluting sectors experience more volatile energy CPI inflation. Moreover, under the ETS, a higher share of free emission allowances allocated to each country tends to reduce inflation volatility, possibly by lowering the effective costs incurred by producers and, thus, their incentives to adjust prices. Finally, our results hold robust to the inclusion of multiple variables of control, alternative identification strategies, different empirical specifications, and subsamples of countries.

These results also have important implications for monetary policy. Given that existing literature has predominantly shown that carbon pricing initiatives tend to have a transitory impact on inflation levels, central banks with credible monetary policies typically overlook this effect, viewing it akin to a negative supply-side shock. However, the fact that ETS schemes could influence inflation volatility introduces a layer of complexity to monetary policy. As jurisdictions seek to expand both the scope and pricing of carbon mechanisms to meet the ambitious climate objectives outlined in the Paris Agreement, regions heavily reliant on ETS schemes may have to deal with heightened inflation volatility. This scenario could complicate the communication and the conduct of monetary policy, presenting central banks with a more intricate signal-extraction challenge. Furthermore, if, as suggested in prior studies (Arndt and Enders 2024), inflation volatility magnifies the transmission of supply shocks to consumer prices, central banks may find it necessary to adapt their strategies to accommodate these additional pressures. Therefore, central banks should closely monitor these effects. Finally, our results contribute to the ongoing debate on the choice of the optimal carbon pricing instrument by shedding light on novel side effects associated with ETS initiatives compared to carbon taxes.

Appendix A. Additional Figures and Tables

Figure A.1. Standard Deviation of CPI Inflation



Source: Authors' own elaboration based on OECD data.

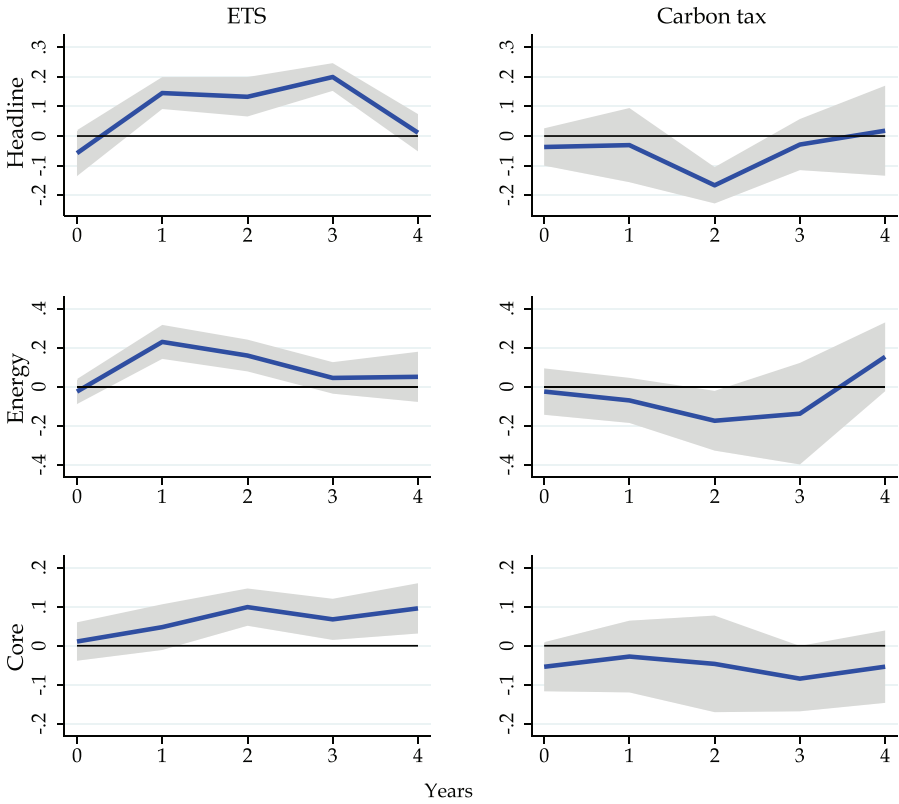
Table A.1. LP: Robustness Tests. Impulse Responses of Headline, Energy, and Core Inflation Volatility for Shocks in Effective ETS Prices

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
	<i>Headline Inflation</i>				
Baseline Estimates	-0.14 (0.20)	0.30 (0.13)	0.21 (0.21)	0.60 (0.20)	0.10 (0.15)
Baseline Period Pre-Paris Agreement	-0.14 (0.25)	0.25 (0.14)	0.21 (0.20)	0.63 (0.34)	0.26 (0.26)
Baseline + OECD Stringency	-0.08 (0.19)	0.29 (0.14)	0.16 (0.23)	0.50 (0.16)	0.08 (0.16)
Baseline + Time FE	-0.12 (0.22)	0.13 (0.12)	0.14 (0.21)	0.69 (0.24)	0.07 (0.16)
	<i>Energy Inflation</i>				
Baseline Estimates	-0.30 (0.39)	1.35 (0.61)	1.15 (0.64)	0.71 (0.61)	0.74 (0.96)
Baseline Period Pre-Paris Agreement	0.12 (0.57)	1.30 (0.83)	1.57 (0.62)	0.74 (0.89)	0.78 (1.21)
Baseline + OECD Stringency	-0.31 (0.45)	1.46 (0.70)	1.16 (0.71)	0.60 (0.59)	0.99 (1.06)
Baseline + Time FE	0.17 (0.36)	0.83 (0.57)	0.71 (0.60)	1.08 (0.65)	0.85 (0.96)
	<i>Core Inflation</i>				
Baseline Estimates	0.04 (0.08)	0.08 (0.13)	0.18 (0.11)	0.13 (0.13)	0.27 (0.19)
Baseline Period Pre-Paris Agreement	0.08 (0.09)	0.06 (0.14)	0.15 (0.15)	0.09 (0.14)	0.27 (0.25)
Baseline + OECD Stringency	0.05 (0.09)	0.15 (0.11)	0.12 (0.09)	0.11 (0.14)	0.19 (0.17)
Baseline + Time FE	0.02 (0.08)	0.10 (0.11)	0.22 (0.14)	0.19 (0.14)	0.34 (0.19)
<p>Note: Impulse responses to an increase in the effective carbon price by USD 40 with 30 percent emission coverage. Difference between point estimates and symmetric 95 percent confidence bands reported in parentheses. All responses are estimated by means of local projections. Baseline estimates include country fixed effects and controlling for GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices: energy inflation, agriculture inflation, and metals and minerals inflation; and CPI inflation. When considering time fixed effects, commodity prices variables are not included. Standard errors are heteroskedasticity robust and clustered by country.</p>					

Table A.2. LP: Robustness Tests. Impulse Responses of Headline, Energy, and Core Inflation Volatility for Shocks in Effective Carbon Tax Rates

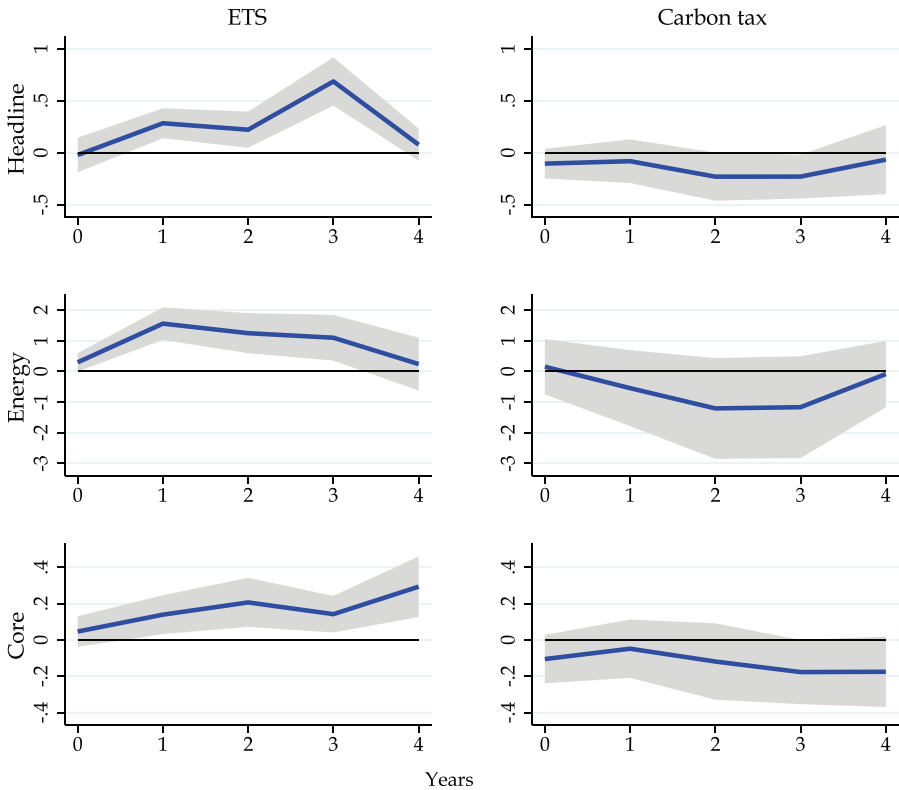
	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
	<i>Headline Inflation</i>				
Baseline Estimates	-0.08 (0.13)	-0.06 (0.27)	-0.39 (0.15)	-0.03 (0.21)	0.06 (0.31)
Baseline Period Pre-Paris Agreement	-0.07 (0.39)	0.12 (0.26)	-0.53 (0.16)	-0.14 (0.30)	-0.07 (0.51)
Baseline + OECD Stringency	-0.08 (0.13)	-0.02 (0.26)	-0.37 (0.17)	-0.06 (0.16)	0.04 (0.33)
Baseline + Time FE	-0.13 (0.18)	0.07 (0.28)	-0.18 (0.26)	-0.09 (0.12)	0.04 (0.32)
	<i>Energy Inflation</i>				
Baseline Estimates	0.19 (0.94)	-0.16 (1.11)	-1.27 (1.11)	-0.50 (1.82)	0.69 (1.13)
Baseline Period Pre-Paris Agreement	0.67 (1.83)	-0.74 (1.28)	-2.58 (1.41)	-2.22 (1.69)	1.07 (1.56)
Baseline + OECD Stringency	0.26 (0.90)	0.06 (1.05)	-1.12 (0.94)	-0.39 (1.63)	0.86 (1.04)
Baseline + Time FE	-0.04 (0.77)	0.22 (1.29)	-0.22 (1.80)	-0.77 (2.10)	0.21 (1.56)
	<i>Core Inflation</i>				
Baseline Estimates	-0.11 (0.12)	-0.06 (0.19)	-0.12 (0.23)	-0.17 (0.15)	-0.09 (0.18)
Baseline Period Pre-Paris Agreement	-0.20 (0.23)	-0.05 (0.20)	-0.14 (0.19)	-0.23 (0.20)	-0.16 (0.24)
Baseline + OECD Stringency	-0.11 (0.12)	-0.07 (0.19)	-0.10 (0.21)	-0.18 (0.13)	-0.14 (0.15)
Baseline + Time FE	-0.11 (0.12)	-0.07 (0.21)	-0.07 (0.22)	-0.12 (0.14)	-0.10 (0.18)
<p>Note: Impulse responses to a USD 40 carbon tax with 30 percent emission coverage. Difference between point estimates and symmetric 95 percent confidence bands reported in parentheses. All responses are estimated by means of local projections. Baseline estimates include country fixed effects and controlling for GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices: energy inflation, agriculture inflation, and metals and minerals inflation; and CPI inflation. When considering time fixed effects, commodity prices variables are not included. Standard errors are heteroskedasticity robust and clustered by country.</p>					

**Figure A.2. Robustness Check:
Results with Log Volatility**



Note: The solid blue lines describe the dynamic impulse responses to an increase in the effective carbon price by USD 40 with 30 percent emission coverage. The gray area delimits 95 percent confidence bands. All responses are estimated by local projections including country fixed effects and controlling for GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices: energy inflation, agriculture inflation, and metals and minerals inflation; and CPI inflation. Standard errors are heteroskedasticity robust and clustered by country.

**Figure A.3. Robustness Check:
Contemporaneous Control Variables**



Note: The solid blue lines describe the dynamic impulse responses to an increase in the effective carbon price by USD 40 with 30 percent emission coverage. The gray area delimits 95 percent confidence bands. All responses are estimated by local projections including country fixed effects and controlling for contemporaneous values of GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices: energy inflation, agriculture inflation, and metals and minerals inflation; and CPI inflation. Standard errors are heteroskedasticity robust and clustered by country.

Table A.3. LP-IV: Kleibergen-Paap Test for Each Horizon and Dependent Variable

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
	<i>Headline Inflation</i>				
Kleibergen and Paap p-value	27.99 (0.00)	25.93 (0.00)	25.76 (0.00)	25.48 (0.00)	25.47 (0.00)
	<i>Energy Inflation</i>				
Kleibergen and Paap p-value	30.38 (0.00)	28.15 (0.00)	27.74 (0.00)	27.24 (0.00)	27.57 (0.00)
	<i>Core Inflation</i>				
Kleibergen and Paap p-value	27.73 (0.00)	25.09 (0.00)	24.73 (0.00)	24.50 (0.00)	24.34 (0.00)
Note: The null hypothesis of the Kleibergen-Paap LM test is that the structural equation is underidentified (i.e., the rank condition fails).					

Table A.4. Classification of Initiatives According to Sectoral Coverage, with Particular Regard to the Electricity Sector

National and Supranational ETSS		
System	Includes Electricity Sector	Other Sectors Covered
EU ETS	Yes	
Switzerland ETS	Yes	
Korea ETS	Yes	
Canada Federal OBPS	No	
Mexico Pilot ETS	Yes	
National Carbon Taxes		
Instrument Name	Includes Electricity Sector	Other Sectors Covered
Poland Carbon Tax		Mining and extractives, Transport, Agriculture, forestry and fishing fuel use
Norway Carbon Tax		Industry, Transport, Aviation, Buildings, Agriculture, forestry and fishing fuel use
Sweden Carbon Tax	Yes	Transport, Buildings, Agriculture, forestry and fishing fuel use
Denmark Carbon Tax	Yes	Transport, Buildings, Agriculture, forestry and fishing fuel use
Latvia Carbon Tax	Yes	Transport, Buildings, Agriculture, forestry and fishing fuel use
Estonia Carbon Tax	Yes	Industry, Buildings
Switzerland Carbon Tax	Yes	Industry, Transport, Buildings, Agriculture, forestry and fishing fuel use
Iceland Carbon Tax	Yes	Industry, Transport, Buildings
Japan Carbon Tax	Yes	Industry, Aviation, Buildings
UK Carbon Price Support	Yes	Transport, Buildings, Agriculture, forestry and fishing fuel use
Mexico Carbon Tax	Yes	Industry
Portugal Carbon Tax	Yes	Industry, Transport, Aviation, Buildings, Agriculture, forestry and fishing fuel use
Chile Carbon Tax	Yes	Industry, Transport, Aviation, Buildings, Agriculture, forestry and fishing fuel use
Colombia Carbon Tax		Transport, Aviation, Buildings, Agriculture, forestry and fishing fuel use
Canada Federal Fuel Charge		Transport, Aviation, Buildings, Agriculture, forestry and fishing fuel use

Source: Carbon Pricing Dashboard. Initiatives classified as including the electricity sector only “in principle” have been excluded from the analysis in order to ensure the accuracy and reliability of our estimations.

Appendix B. Panel Data Analysis

Empirical Specification

An alternative panel data specification can also be outlined as follows:

$$\sigma(\pi)_{i,t} = \alpha\sigma(\pi)_{i,t-1} + \beta\textit{Carbon pricing}_{i,t} + \gamma X_{i,t-1} + \mu_i + \epsilon_{i,t}, \quad (\text{B.1})$$

where $\sigma(\pi)_{i,t}$ is the dependent variable measured, for a country i and for each year t , as the standard deviation of monthly year-on-year inflation computed in windows of 24 months.

Carbon pricing $_{i,t}$ includes the set of variables that depict the existence and different features of carbon pricing systems, representing our variables of interest. *Carbon pricing, any* is a dummy variable that takes value 1 if the country has any type—national, subnational, or regional—of carbon pricing implemented for a given year, while *ETS, any* and *Carbon tax, any* reflect whether the country has implemented any type of ETS or carbon tax, respectively. *Coverage-fully overlapping* and *Coverage-not overlapping* measure the amount of GHG emissions covered by the initiatives in that country. $X_{i,t-1}$ includes the lag of the relevant controls described in Section 3. Finally, μ_i is a country-level time-invariant effect that reflects individual heterogeneity, and $\epsilon_{i,t}$ represents an idiosyncratic disturbance that is expected to be independent and identically distributed. β and γ are the coefficients associated with interest and control variables, respectively.

However, if the lag of the dependent variable is not significant the above econometric specification turns into the following static panel:

$$\sigma(\pi)_{i,t} = \beta\textit{Carbon pricing}_{i,t} + \gamma X_{i,t-1} + \mu_i + \epsilon_{i,t}, \quad (\text{B.2})$$

which could be estimated through the fixed-effects or random-effects estimators.²⁵

²⁵In the presence of correlation between the individual effect and the explanatory variables, the fixed-effects estimator should be employed, as the random-effects estimator is expected to yield biased and inconsistent estimates of the regression parameters. The Hausman test is usually employed to opt for the

Panel Data Results

Our baseline results for panel data are presented in Tables B.1 and B.2. Given that the lag of the inflation volatility is found not significant,²⁶ we focus on the results of the static panel.^{27,28}

Table B.1 depicts a series of specifications with the objective to test the effects of the existence of any carbon pricing initiative, without further differencing by types. Columns 1–5 present some stepwise estimations in which the relevant groups of control variables are sequentially introduced in the specification of the models. In a first step, general socioeconomic variables that account for the size of the country and the level of development are included (column 1). Afterward, trade openness (column 2), controls for exposure to commodity prices and capital flows (column 3), variables reflecting common global factors, mainly commodity prices (column 4), and, finally, the domestic level of CPI inflation (column 5) are added. Table B.2 presents the same structure as Table B.1, but in this case carbon pricing systems are further disaggregated as to consider the ETS and carbon tax types separately.

Overall, the existence of a carbon pricing systems is related to increased volatility of inflation (Table B.1), which seems to be solely driven by ETS schemes (Table B.2). Carbon pricing systems in the form of a carbon tax do not seem to have a significant impact on inflation volatility in any of the specifications (Table B.2).

In a second step, we investigate the channels by which carbon pricing systems affect inflation volatility, analyzing the differential impact on the energy CPI component and the degree of pass-through

fixed- or the random-effects estimator. The random-effects estimator is preferred under the null hypothesis (zero correlation between the individual effect and the regressors) due to higher efficiency, while under the alternative the fixed-effects estimator is at least consistent and, thus, preferred.

²⁶Estimates of the System-GMM estimator differentiating by type of carbon pricing systems are available upon request. In none of the specifications considered is the lag of inflation volatility significant.

²⁷For all the models, results for the Hausman test reject the null hypothesis that individual effects are uncorrelated with the regressors. Therefore, we present the results of the fixed-effects estimator.

²⁸We also performed a test for normality of the error components in panel data models proposed by Alejo et al. (2015). Our post-estimation results show that the null hypothesis of joint normality of both error terms (individual effect and idiosyncratic component) cannot be rejected.

Table B.1. Baseline Results: Carbon Pricing

	(1)	(2)	(3)	(4)	(5)
Carbon Pricing, Any	0.419*** (0.000)	0.450*** (0.000)	0.464*** (0.000)	0.512*** (0.000)	0.343*** (0.000)
GDP per Capita	-0.047*** (0.000)	-0.038*** (0.001)	-0.044*** (0.000)	-0.045*** (0.000)	-0.029*** (0.002)
Population	-0.048*** (0.000)	-0.050*** (0.000)	-0.062*** (0.000)	-0.060*** (0.000)	-0.021** (0.035)
Trade Openness		-0.004 (0.169)	-0.005 (0.110)	-0.005* (0.085)	-0.002 (0.326)
Primary Exporter			0.279* (0.088)	0.249 (0.114)	0.145 (0.239)
Capital Flows			-0.000 (0.744)	0.000 (0.997)	-0.000 (0.598)
Commodity Prices: Energy Inflation				0.003 (0.154)	-0.002 (0.199)
Commodity Prices: Agriculture Inflation				0.024*** (0.000)	0.022*** (0.000)
Commodity Prices: Metals and Minerals Inflation				-0.013*** (0.000)	-0.011*** (0.000)
CPI Inflation					0.164*** (0.000)
Constant	4.204*** (0.000)	4.321*** (0.000)	4.983*** (0.000)	4.959*** (0.000)	2.397*** (0.000)
Observations	756	756	708	708	708
R-squared	0.059	0.061	0.078	0.142	0.477
Number of Countries	36	36	36	36	36
Country FE	Yes	Yes	Yes	Yes	Yes

Note: *, **, and *** denote statistical significance at 10 percent, 5 percent, and 1 percent, respectively. Robust p-values are in parentheses. See main text of the paper for the definition of the variables.

Table B.2. Baseline Results: ETSs and Carbon Taxes

	(1)	(2)	(3)	(4)	(5)
ETS, Any	0.400*** (0.001)	0.429*** (0.000)	0.426*** (0.001)	0.464*** (0.000)	0.393*** (0.000)
Carbon Tax, Any	0.026 (0.850)	0.032 (0.816)	0.031 (0.829)	0.095 (0.504)	0.026 (0.813)
GDP per Capita	-0.052*** (0.000)	-0.044*** (0.000)	-0.050*** (0.000)	-0.051*** (0.000)	-0.036*** (0.000)
Population	-0.041*** (0.000)	-0.042*** (0.000)	-0.052*** (0.000)	-0.050*** (0.000)	-0.014 (0.142)
Trade Openness		-0.004 (0.169)	-0.005 (0.113)	-0.005* (0.081)	-0.003 (0.251)
Primary Exporter			0.211 (0.202)	0.175 (0.274)	0.079 (0.527)
Capital Flows			-0.000 (0.732)	0.000 (0.989)	-0.000 (0.559)
Commodity Prices: Energy Inflation				0.004* (0.087)	-0.001 (0.398)
Commodity Prices: Agriculture Inflation				0.023*** (0.000)	0.021*** (0.000)
Commodity Prices: Metals and Minerals Inflation				-0.013*** (0.000)	-0.011*** (0.000)
CPI Inflation					0.165*** (0.000)
Constant	4.197*** (0.000)	4.319*** (0.000)	4.925*** (0.000)	4.909*** (0.000)	2.449*** (0.000)
Observations	756	756	708	708	708
R-squared	0.058	0.060	0.076	0.140	0.480
Number of Countries	36	36	36	36	36
Country FE	Yes	Yes	Yes	Yes	Yes

Note: *, **, and *** denote statistical significance at 10 percent, 5 percent, and 1 percent, respectively. Robust p-values are in parentheses. See main text of the paper for the definition of the variables.

Table B.3. Energy and Core Inflation

	(1) Sd. Headline	(2) Sd. Energy Infl.	(3) Sd. Core	(4) Sd. Headline	(5) Sd. Energy	(6) Sd. Core
Carbon Pricing, Any	0.343*** (0.000)	0.661* (0.063)	0.337*** (0.000)	0.393*** (0.000)	1.258*** (0.001)	0.415*** (0.000)
ETS, Any				0.026 (0.813)	-0.082 (0.839)	0.145 (0.150)
Carbon Tax, Any						
Observations	708	708	707	708	708	707
R-squared	0.477	0.359	0.377	0.480	0.368	0.388
Number of Countries	36	36	36	36	36	36
Control Variables Included	Yes	Yes	Yes	Yes	Yes	Yes

Note: ***, **, and * denote significance at 1 percent, 5 percent, and 10 percent, respectively. Robust p-values in parentheses. Control variables: GDP per capita; population; trade openness; primary exporter; capital flows; commodity prices; energy inflation, agriculture inflation, and metals and minerals inflation; CPI inflation. See main text of the paper for the definition of the variables.

to the volatility of core inflation. For this purpose, in Table B.3 the dependent variables are headline, energy, and core inflation volatility, respectively. We consider the overall effect of having implemented any carbon pricing initiative (columns 1, 2, and 3) and the effect of ETSs and carbon taxes separately (columns 4, 5, and 6). Control variables from the baseline specification are also included, in a similar vein as in column 5 in Tables B.1 and B.2.

As in our baseline local projections specification, the estimates reveal that carbon pricing systems do significantly impact inflation volatility for the case of energy CPI, which should be expected given that energy-related goods are the ones that are targeted directly or indirectly by carbon pricing initiatives. Regarding the type of carbon pricing initiative affecting inflation volatility, again, it seems that schemes based on the ETS are the ones that spur the increased volatility pattern, affecting headline, energy, and core inflation volatility (columns 5 and 6). We cannot find a significant effect of carbon taxes on the volatility of any constituent of the CPI.

References

- Adolfson, J. F., M. F. Minnesso, J. E. Mork, and I. V. Robays. 2024. "Gas Price Shocks and Euro Area Inflation." ECB Working Paper No. 2905.
- Aghion, P., P. Bacchetta, and A. Banerjee. 2004. "Financial Development and the Instability of Open Economies." *Journal of Monetary Economics* 51 (6): 1077–1106.
- Alejo, J., A. Galvao, G. Montes-Rojas, and W. Sosa-Escudero. 2015. "Tests for Normality in Linear Panel-Data Models." *Stata Journal* 15 (3): 822–32.
- Andersson, J. J. 2019. "Carbon Taxes and CO₂ Emissions: Sweden as a Case Study." *American Economic Journal: Economic Policy* 11 (4): 1–30.
- Annicchiarico, B., S. Carattini, C. Fischer, and G. Heutel. 2021. "Business Cycles and Environmental Policy: Literature Review and Policy Implications." NBER Working Paper No. 29032.
- Annicchiarico, B., and F. Di Dio. 2017. "GHG Emissions Control and Monetary Policy." *Environmental and Resource Economics* 67 (4): 823–51.

- Arndt, S., and Z. Enders. 2024. "The Transmission of Supply Shocks in Different Inflation Regimes." Working Paper No. 938, Banque de France.
- Balatti, M. 2020. "Inflation Volatility in Small and Large Advanced Open Economies." ECB Working Paper No. 2448.
- Batten, S., R. Sowerbutts, and M. Tanaka. 2020. "Climate Change: Macroeconomic Impact and Implications for Monetary Policy." In *Ecological, Societal, and Technological Risks and the Financial Sector*, ed. T. Walker, D. Gramlich, M. Bitar, and P. Fardnia, 13–38. Palgrave Macmillan.
- Blanchard, O., E. Cerutti, and L. H. Summers. 2015. "Inflation and Activity: Two Explorations and Their Monetary Policy Implications." Working Paper No. 15-19, Peterson Institute for International Economics.
- Bleaney, M., and D. Fielding. 2002. "Exchange Rate Regimes, Inflation and Output Volatility in Developing Countries." *Journal of Development Economics* 68 (1): 233–45.
- Bowdler, C., and A. Malik. 2005. "Openness and Inflation Volatility: Cross-Country Evidence." Economics Paper No. 2005-W14, Economics Group, Nuffield College, University of Oxford.
- . 2017. "Openness and Inflation Volatility: Panel Data Evidence." *North American Journal of Economics and Finance* 41 (July): 57–69.
- Byrne, J. P., and E. Davis. 2004. "Permanent and Temporary Inflation Uncertainty and Investment in the United States." *Economics Letters* 85 (2): 271–77.
- Carattini, S., G. Heutel, and G. Melkadze. 2021. "Climate Policy, Financial Frictions, and Transition Risk." NBER Working Paper No. 28525.
- Cukierman, A., and A. H. Meltzer. 1986. "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information." *Econometrica* 54 (5): 1099–1128.
- Del Negro, M., J. di Giovanni, and K. Dogra. 2023. "Is the Green Transition Inflationary?" Staff Report No. 1053, Federal Reserve Bank of New York.
- Diluiso, F., B. Annicchiarico, M. Kalkuhl, and J. C. Minx. 2021. "Climate Actions and Macro-financial Stability: The Role of Central Banks." *Journal of Environmental Economics and Management* 110 (October): Article 102548.

- Economides, G., and A. Xepapadeas. 2018. "Monetary Policy under Climate Change." CESifo Working Paper No. 7021.
- Elder, J. 2004. "Another Perspective on the Effects of Inflation Uncertainty." *Journal of Money, Credit and Banking* 36 (5): 911–28.
- Elkins, P., and T. Baker. 2001. "Carbon Taxes and Carbon Emissions Trading." *Journal of Economic Surveys* 15 (3): 325–76.
- European Central Bank. 2021. "Climate Change and Monetary Policy in the Euro Area." Occasional Paper No. 271.
- Fernández, A., S. Schmitt-Grohé, and M. Uribe. 2017. "World Shocks, World Prices, and Business Cycles: An Empirical Investigation." *Journal of International Economics* 108 (Supplement 1): S2–S14.
- Ferrari, A., and V. Nispi Landi. 2023. "Toward a Green Economy: The Role of the Central Bank's Asset Purchases." *International Journal of Central Banking* 19 (5): 287–340.
- Forbes, K. 2019. "Inflation Dynamics: Dead, Dormant, or Determined Abroad?" NBER Working Paper No. 26496.
- Friedman, M. 1977. "Nobel Lecture: Inflation and Unemployment." *Journal of Political Economy* 85 (3): 451–72.
- Gertler, M., and P. Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7 (1): 44–76.
- Goulder, L. H., and A. R. Schein. 2013. "Carbon Taxes versus Cap and Trade: A Critical Review." *Climate Change Economics* 04 (03): 1–28.
- 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.
- Ha, J., A. Kose, and F. Ohnsorge, eds. 2019. *Inflation in Emerging and Developing Economies: Evolution, Drivers and Policies*. World Bank.
- Hepburn, C., J. Stiglitz, and N. Stern. 2020. "Carbon Pricing." *European Economic Review* 127 (August): Article 103440.
- Intergovernmental Panel on Climate Change (IPCC). 2023. "Climate Change 2023: Synthesis Report." Contribution of Working Groups I, II and III to the Sixth Assessment Report of

- the Intergovernmental Panel on Climate Change, IPCC, Geneva, Switzerland.
- Jordà, Ò. 2005. “Estimation and Inference of Impulse Responses by Local Projections.” *American Economic Review* 95 (1): 161–82.
- Jordà, Ò., M. Schularick, and A. M. Taylor. 2015. “Betting the House.” *Journal of International Economics* 96 (Supplement 1): S2–S18.
- Judson, R., and A. Orphanides. 1999. “Inflation, Volatility and Growth.” *International Finance* 2 (1): 117–38.
- Kaldorf, M., and F. Giovanardi. 2023. “Climate Change and the Macroeconomics of Bank Capital Regulation.” Working Paper.
- Kamber, G., and B. Wong. 2020. “Global Factors and Trend Inflation.” *Journal of International Economics* 122 (January): Article 103265.
- Känzig, D. R. 2021. “The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements.” *American Economic Review* 111 (4): 1092–1125.
- . 2023. “The Unequal Economic Consequences of Carbon Pricing.” NBER Working Paper No. 31221.
- Känzig, D. R., and M. Konradt. 2023. “Climate Policy and the Economy: Evidence from Europe’s Carbon Pricing Initiatives.” NBER Working Paper No. 31260.
- Kilian, L., and X. Zhou. 2023. “A Broader Perspective on the Inflationary Effects of Energy Price Shocks.” *Energy Economics* 125 (September): Article 106893.
- Kleibergen, F., and P. Paap. 2006. “Generalized Reduced Rank Tests Using the Singular Value Decomposition.” *Journal of Econometrics* 133 (1): 97–126.
- Konradt, M., and B. Weder di Mauro. 2021. “Carbon Taxation and Greenflation: Evidence from Europe and Canada.” IHEID Working Paper No. 17-2021, Economics Section, Graduate Institute of International Studies.
- . 2023. “Carbon Taxation and Greenflation: Evidence from Europe and Canada.” *Journal of the European Economic Association* 21 (6): 2518–46.
- Kuttner, K. N. 2001. “Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market.” *Journal of Monetary Economics* 47 (3): 523–44.

- Lane, P. R. 1997. "Inflation in Open Economies." *Journal of International Economics* 42 (3–4): 327–47.
- López, L., S. Párraga Rodríguez, and D. Santabárbara. 2022. "Box 4. The Pass-through of Higher Natural Gas Prices to Inflation in the Euro Area and in Spain." *Economic Bulletin* (Banco de España) 3/2022: 49–52.
- Martin, R., M. Muûls, L. B. de Preux, and U. J. Wagner. 2014. "Industry Compensation under Relocation Risk: A Firm-Level Analysis of the EU Emissions Trading Scheme." *American Economic Review* 104 (8): 2482–2508.
- Martin, R., M. Muûls, and U. J. Wagner. 2016. "The Impact of the European Union Emissions Trading Scheme on Regulated Firms: What is the Evidence after Ten Years?" *Review of Environmental Economics and Policy* 10 (1): 129–48.
- McKibbin, W., M. Konradt, and B. Weder di Mauro. 2021. "Climate Policies and Monetary Policies in the Euro Area." Paper for the ECB Sintra Forum 2021.
- McKibbin, W., A. Morris, and P. Wilcoxon. 2014. "The Economic Consequences of Delay in US Climate Policy." CAMA Working Paper No. 2014-49.
- Metcalf, G. E. 2021. "Carbon Taxes in Theory and Practice." *Annual Review of Resource Economics* 13: 245–65.
- Metcalf, G. E., and J. H. Stock. 2020. "Measuring the Macroeconomic Impact of Carbon Taxes." *AEA Papers and Proceedings* 110 (May): 101–6.
- . 2023. "The Macroeconomic Impact of Europe's Carbon Taxes." *American Economic Journal: Macroeconomics* 15 (3): 265–86.
- Moessner, R. 2022. "Effects of Carbon Pricing on Inflation." CESifo Working Paper No. 9563.
- Mundaca, G., J. Strand, and I. R. Young. 2021. "Carbon Pricing of International Transport Fuels: Impacts on Carbon Emissions and Trade Activity." *Journal of Environmental Economics and Management* 110 (October): Article 102517.
- Nakamura, E., and J. Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics* 133 (3): 1283–1330.
- Nakov, A., and C. Thomas. 2023. "Climate-Conscious Monetary Policy." Working Paper No. 2334, Banco de España.

- Neely, C. J., and D. E. Rapach. 2011. "International Comovements in Inflation Rates and Country Characteristics." *Journal of International Money and Finance* 30 (7): 1471–90.
- Network for Greening the Financial System (NGFS). 2021. "Annual Report 2020."
- Parker, M. 2018. "How Global is 'Global Inflation'?" *Journal of Macroeconomics* 58 (December): 174–97.
- Parry, I., S. Black, and K. Zhunussova. 2022. "Carbon Taxes or Emissions Trading Systems?: Instrument Choice and Design." Staff Climate Note No. 2022/006, International Monetary Fund.
- Pigou, A. 1933. "The Economics of Welfare. by A. Pigou." (Review of Fourth Edition.) *Economic Journal* 43 (170): 329–30.
- Romer, D. 1993. "Openness and Inflation: Theory and Evidence." *Quarterly Journal of Economics* 108 (4): 869–903.
- Schnabel, I. 2022. "Looking Through Higher Energy Prices? Monetary Policy and the Green Transition." Speech at a panel on "Climate and the Financial System" at the American Financial Association 2022 Virtual Annual Meeting, Frankfurt am Main, January 8.
- Shapiro, A., and G. E. Metcalf. 2021. "The Macroeconomic Effects of a Carbon Tax to Meet the U.S. Paris Agreement Target: The Role of Firm Creation and Technology Adoption." NBER Working Paper No. 28795.
- 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.
- Venmans, F., J. Ellis, and D. Nachtigall. 2020. "Carbon Pricing and Competitiveness: Are They at Odds?" *Climate Policy* 20 (9): 1070–91.
- Weitzman, M. L. 1974. "Prices vs. Quantities." *Review of Economic Studies* 41 (4): 477–91.
- World Bank. 2021. *State and Trends of Carbon Pricing 2021*. Washington, DC: World Bank.
- Yamazaki, A. 2017. "Jobs and Climate Policy: Evidence from British Columbia's Revenue-Neutral Carbon Tax." *Journal of Environmental Economics and Management* 83 (May): 197–216.

Asymmetric Shocks and Monetary Policy in the Euro Area*

Davor Kunovac,^{a,c} Diego Rodriguez Palenzuela,^b
and Yiqiao Sun^b

^aHrvatska narodna banka

^bEuropean Central Bank

^cUniversity of Rijeka, Faculty of Economics and Business

We analyze the role of asymmetric macroeconomic shocks for conduct of monetary policy of the euro area, as a monetary union. Taking an international business cycle perspective, we distinguish between symmetric shocks, which affect countries similarly, and asymmetric shocks, which cause diverging movements across countries. Applying a Bayesian sign- and zero-restricted structural vector autoregression (BVAR) model, we find that global symmetric shocks predominantly drive business cycles in most euro-area countries. Linked to this, we find evidence of a euro-area-specific cycle. Through the lens of our framework, we also find empirical support for the optimum currency area (OCA) endogeneity hypothesis: upon adopting the euro, new member countries, on average, experience faster convergence in their share of symmetric shocks relative to the euro-area aggregate. Our econometric framework also allows us to derive an optimum currency area index for the euro area based on a *shocks-similarity principle*. Our index captures not only the relative importance of symmetric shocks but also the extent of cross-country heterogeneity. This OCA index highlights that, while the OCA features of the euro area remain

*We thank an anonymous referee, Ivan Zilic, Lucija Fioretti, Karlo Kotarac, Jurica Zrnc, and Milan Deskar-Škrbić for valuable comments and helpful suggestions. We also thank, for their comments, participants at the International Conference “Reassessment of the Optimal Currency Area Theory in the Persistently Heterogeneous European Union” in Vilnius. The views expressed are those of the authors only and do not necessarily reflect those of the Eurosystem, the European Central Bank, or the Hrvatska narodna banka. A version of this paper was previously circulated under the title “A New Optimum Currency Area Index for the Euro Area.” Author e-mails: davor.kunovac@hnb.hr (corresponding author), diego.rodriquez@ecb.europa.eu, yiqiao.sun@ecb.europa.eu.

firmly established, their further improvement has stagnated over time. All in all, current levels of the OCA index suggest a broadly similar level as at the start of the euro, underscoring the importance of not being complacent with the policy goal of EU single market integration.

JEL Codes: F33, F44, E42, E52.

1. Introduction

It is well known that when a set of countries or regions forgo their sovereign currency and independent monetary policy to form a shared monetary union, important gains in terms of economic integration may be attained, as exchange-rate-related transaction costs and volatility are eliminated and price transparency is enhanced. To assess the net gains in forming the union, benefits need to be compared to the costs of losing monetary autonomy at country level: Once in the monetary union, the country has lost the full power of an important instrument of macroeconomic stabilization, as monetary policy has demonstrated to be a key policy tool to address macroeconomic shocks, in particular of a demand nature. The costs of lost monetary independence will be attenuated, however, if the various countries or regions forming the union are subject to broadly similar or symmetric macroeconomic shocks. In the limit, if macroeconomic shocks were perfectly correlated across countries, the cost of foregone monetary independence would be nihil. Understanding the relative importance of symmetric versus asymmetric macroeconomic shocks is thus crucial for assessing the costs of forming a monetary union, and ultimately its desirability: When the business cycles of the union members, at country or regional level, are primarily driven by common or symmetric shocks and, therefore, positively correlated, the single *one-size-fits-all* monetary policy is in a stronger position to stabilize all countries in the union simultaneously; see Alesina, Barro, and Tenreyro (2002). Conversely, asymmetric shocks, with an uneven impact across countries, tend to impair the conduct of the common monetary policy.

These same premises determine the net benefits for a single country in *joining* an established and ongoing monetary union, as is the case for candidate countries considering the best timing for joining the euro area, such as the Czech Republic and Romania are

currently. If the candidate countries' business cycles are sufficiently synchronized with those of the euro area, the costs of giving up their sovereign currency and monetary policy should be of lesser concern.

The role of asymmetric shocks has been from the outset at the center of optimum currency area (OCA) theory, a widely accepted framework for assessing whether the conditions for a group of countries to form a currency union are met. At the initial development of the OCA theory, seminal contributions from Mundell (1961), McKinnon (1963), and Kenen (1969) identified a number of prerequisites, in the form of necessary structural conditions, for successfully forming a currency union. Prominent among these were nominal flexibility in goods and labor markets, factor mobility, production diversification, and fiscal integration. In a subsequent phase of OCA theory, the focus shifted to the costs of fixing irreversible exchange rates and adopting a common currency, along with the role of financial integration; see Corden (1972), Ishiyama (1975), and Mundell (1975).

The combination of the former conditions and latter cost features gave rise to a set of *OCA criteria*, that is, a set of indicators to assess the viability and desirability of forming a monetary union. Yet, while the identified criteria are conceptually well grounded, in practice it has proven difficult to measure empirically the degree to which some of the criteria are fulfilled, both *ex ante* or on an ongoing basis. Much of the empirical literature on OCA conditions, such as Bayoumi and Eichengreen (1992), Masson and Taylor (1993), Alesina, Barro, and Tenreyro (2002), and Mongelli (2002), suggested that *similarity of shocks* driving union members would qualify as a *catch-all* OCA property capturing the interaction between several of these properties at once. However, one important aspect, the measurement of the similarity of macro shocks across countries, has remained a particularly daunting task.

To better understand the similarity of shocks and business cycle coherence across members, the existing literature, strongly influenced by Bayoumi and Eichengreen (1992), mostly weighs the cross-country correlation of the measurable proxies for shocks. This is done usually without placing emphasis on the separate identification of symmetric versus asymmetric shocks. While this approach is intuitive and easily implementable, it clearly falls short of capturing the structural diversity of shocks and the complexity of shock

transmission mechanisms across diverse economies. This important gap in the literature is precisely where our paper makes a contribution. Instead, we take a fundamentally different approach, setting up an econometric framework that permits the identification and measurement of symmetric versus asymmetric shocks and, among the latter, of supply versus demand shocks. Our econometric framework builds on a sign- and zero-restricted Bayesian structural vector autoregression (BVAR) model, allowing for a more structural and reliable analysis of shock dynamics at various levels of aggregation, namely global, euro area, and local (country level).

A further central theme of OCA theory where shock similarity plays a key role relates to how *fit* or well prepared a candidate country stands as a new member of a possibly complex and heterogeneous monetary union. Being well prepared for membership is closely linked to how the business cycle coherence between the candidate country and the monetary union evolves once the country joins the union. The *OCA endogeneity hypothesis*, as proposed by Frankel and Rose (1997, 1998), suggests that business cycles and shock similarities across new members and the rest of the union are not independent. Similarities are, under this view, expected to gradually increase *ex post* and as a consequence of the currency area enlargement: joining the currency union should foster stronger trade linkages and deeper investment connections. Nonetheless, counterarguments to the OCA endogeneity theory have been made. Krugman (1993) suggests that a country joining an existing monetary union could lead to increased specialization in specific sectors of the enlarged union, thereby potentially reducing business cycle coherence, with a deterioration of the union's OCA properties. These conflicting theories highlight the need not only to evaluate the *ex ante* conditions for currency union membership but also to monitor the evolution of *ex post* convergence (in line with the OCA endogeneity view) or lack thereof (as suggested by the increasing sectoral specialization view).

The main contribution of our paper is that we formalize a precise and measurable definition of asymmetric shocks for the euro area, both from a country-specific and a union-wide perspective.¹

¹Although the European Central Bank (ECB) makes it clear the focus of its monetary policy is the euro area as a whole, the country perspective remains

To quantify the relative importance of symmetric versus asymmetric shocks for euro-area members, we rely on a set of identified BVAR models, which yields a shock accounting exercise of interest on its own. This exercise in turn serves as the building block to study in more detail the role and relative importance of symmetric versus asymmetric shocks for the smooth monetary policy conduct and for the suitability of euro adoption for new member countries, as well as the performance of the latter once in the monetary union. These results provide a novel approach and test of the OCA endogeneity hypothesis through the lens of our structural BVAR framework. Finally, as a by-product of our shock accounting exercise, we propose a new OCA index based on the relative importance of symmetric shocks.

To preview our results, we show that symmetric shocks emerge as the dominant drivers of GDP growth across euro-area countries. Among the symmetric shocks, an important fraction are global in nature, which partly explains why the ECB's monetary policy is often well aligned with that of other central banks, notably the U.S. Federal Reserve. Additionally, we find evidence of a significant comovement across euro-area members that is independent from the identified global commonality—i.e., a distinct euro-area business cycle. Our results contribute to ongoing debates in the literature on this topic: while some studies, such as those by Artis (2003) or Camacho, Perez-Quiros, and Saiz (2006), have struggled to identify a clear European cycle, our findings align with those of Inklaar and de Haan (2001), Kose, Otrok, and Whiteman (2003), or Ferroni and Klaus (2015), among others, who argue in favor of its existence and relevance.

While we show that a large proportion of shocks are symmetric, this is not sufficient per se to conclude on a benign assessment on the conditions for the common policy to stabilize the euro area. The effectiveness of the common monetary policy depends strongly on the structural nature and the composition of shocks underlying business cycle dynamics. In particular, a larger prevalence of symmetric *demand* shocks is often related to more favorable conditions

important. In a case in which a member country is affected by different shocks compared to the rest of the union, the common monetary policy cannot be equally desirable for all, which could create tensions between countries or even a possible breakup.

to stabilize the euro area, while, on the other hand, symmetric supply shocks may create a dilemma for central bankers, who may or may not be in a position to react to them, as highlighted by Beaudry, Carter, and Lahiri (2022), Madeira, Madeira, and Monteiro (2023), or Tenreyro et al. (2023). We show that both demand- and supply-side shocks at euro-area level are important in driving the euro-area business cycle dynamic.

While the results highlighted so far bode well for the OCA properties of the euro area, it remains to be checked whether country-level shocks of demand nature are not dominant. Indeed, from a country perspective, of either a candidate new-member country or an existing member, the loss of the independent monetary policy is more severe when idiosyncratic demand shocks are more prevalent. Reassuringly for the assessment of euro-area OCA conditions, we find that country-level shocks from the demand side are relatively rare, suggesting that the costs of lost monetary policy at country level are likely limited. In particular, we find that business cycle dynamics of Central and Eastern European countries—both euro-area members and candidate countries—are becoming increasingly similar to the euro-area aggregate, in terms of shock composition. This finding would seem to mitigate concerns sometimes expressed that euro-area enlargement would increase the risks of asymmetric shocks, as raised, e.g., by De Grauwe (2003). Our results suggest that these risks have diminished over time.

More generally, our framework serves to shed light on the role of euro adoption in the alignment process for business cycle dynamics for countries adopting the euro. For this purpose, we test how the euro adoption affects the importance of symmetric shocks for *new* members, i.e., the countries in our sample that joined at later enlargement stages—Slovakia, Slovenia, Estonia, Latvia, and Lithuania. This is the context to which we apply and test the *OCA endogeneity hypothesis* of Frankel and Rose (1998). To investigate if the euro adoption has indeed promoted business cycle synchronization, we rely on difference-in-difference (D-I-D) estimation. Recent findings in the methodological literature, however, point out drawbacks of a standard D-I-D approach,² whenever there are multiple

²Here, a panel regression with time and country fixed effects and euro membership being a *treatment* dummy.

treatment periods (i.e., different entry dates); see Callaway and Sant'Anna (2021) for an explanation. Therefore, in our case, given that countries enter the euro area in different years, we apply both the standard D-I-D and the method by Callaway and Sant'Anna (2021). The results provide evidence of the positive impact of euro adoption on shock synchronization. This finding may be relevant with respect to the ongoing policy debate on some non-euro-area EU countries considering the euro-area membership.³

Finally, based on the basic *signal-to-noise ratio* concept, we propose an indicator aimed at summarizing our findings on the behavior of both cross-country average and dispersion of relative importance of symmetric shocks. We regard this indicator as a suitable *OCA index*, and we show how it applies for the euro area over time. Our indicator underscores that high relative importance of symmetric shocks across countries alone may not be sufficient to improve OCA features of a currency union, whenever there are important differences between countries. In particular, a monetary union with a high cross-country average importance of symmetric shocks may be characterized as poor in terms of its OCA features in a case where the relative importance of symmetric shocks is overly dispersed across countries. Our OCA index for the euro area thus reflects the view that cross-country heterogeneity in business cycle fluctuations makes the common currency not equally desirable for all area members, which could create tensions between countries (De Grauwe 1996) and eventually even threaten the political viability of EMU (Orphanides 2020). Increased heterogeneity across union members has important implications for policy as well, as the policy stance in a currency union during the crises may be difficult to calibrate, despite given predominance of symmetric shocks on average. Signal-to-noise ratio may, therefore, serve as a useful complement to average importance of symmetric shocks in tracking OCA properties, especially during the crises. Our OCA index is slow-moving and a good reflection of changing underlying economic structures across the euro area and, therefore, informative about the ability of monetary policy to

³For largely positive sentiment toward the euro adoption in candidate countries, see <https://www.brusselstimes.com/858918/czechia-to-take-concrete-steps-to-adopt-the-euro> for Czechia and <https://www.euractiv.com/section/politics/news/romania-wants-to-push-euro-adoption-by-2026/> for Romania.

stabilize the euro-area economy in the medium run. We show that an index that takes both moments into account is able to properly signal different types of crises in a currency union: those that originate in the rest of the world, outside the euro area (e.g., the Global Financial Crisis, or GFC), but also those locally or inside the euro area, originating in one or few euro-area member countries (e.g., the euro debt crisis). Otherwise, an index based on the relative importance of symmetric shocks only remains blind to the latter type of crisis. Through the lens of our OCA indicators, we observe that cyclical convergence in the euro area proceeds gradually, with OCA features showing signs of stagnation at levels close to those seen before the launch of the euro.

2. Relation to the Literature

This paper contributes to several strands of literature on the desirability and performance of forming and over time enlarging a monetary union: first, in connection to the identification and measurement of symmetric and asymmetric shocks among euro-area members and their implications for the common monetary policy; second, on the testing the OCA endogeneity theory; and, third, to the literature on OCA indices.⁴

A symmetric macroeconomic shock refers to an unexpected event that affects all union members similarly and simultaneously (European Commission 1990). Surprisingly, given the importance for its OCA conditions, the differentiation between symmetric and asymmetric shocks affecting euro-area countries has not been a primary focus in the literature. Instead, most studies have emphasized *the similarity of shocks*, a concept initially developed within a structural VAR framework by Bayoumi and Eichengreen (1993). This framework has become widely used, appearing, e.g., in Fidrmuc and Korhonen (2003) and Campos and Macchiarelli (2016), among others. These studies separately identify demand and supply shocks

⁴More broadly, our paper is also naturally related to the literature on international business cycles and, more specifically, to business cycle synchronization among euro-area countries; see, for example, Camacho, Perez-Quiros, and Saiz (2006), Giannone, Lenza, and Reichlin (2008), or Ferroni and Klaus (2015).

across European countries using the Blanchard-Quah methodology (Blanchard and Quah 1989) and assess the costs of forgoing autonomous monetary policy, relying on cross-country shock correlations. However, simple correlations between shocks falls short as an indicator to evaluate the costs of having a common monetary policy. In particular, the correlation between shocks may be high, but their overall macroeconomic importance across countries may be very different. In that case, a common one-size-fits-all monetary policy cannot be equally suitable for all countries in a monetary union and the costs of having a common currency are unevenly distributed across constituent countries. For this reason, we take a different track and distinguish the impact of symmetric from asymmetric shocks using a historical shock decomposition based on BVAR models. In contrast to simple correlation between shocks, a measure based on our shock accounting approach is dynamic in that it takes into account how the shocks propagate through the economy and contribute to overall macroeconomic variability.

The assessment of symmetric shocks has been prevalent in the VAR literature, but mainly in the context of analyzing bilateral linkages between the euro area and non-euro-area countries. Examples include Peersman (2011) for the U.K., Audzei and Brazdik (2018) for selected Central and Eastern European countries, and, more recently, Deskar-Škrbić, Kotarac, and Kunovac (2020) for Croatia, Bulgaria, and Romania, the three candidate countries intending to join the euro area, and Deskar-Škrbić and Kunovac (2020) for EU members (Sweden, Czech Republic, Hungary, and Poland) not intending or ready to join the euro area. The main focus in this paper, in contrast, is on the coherence of business cycles of existing euro-area members with respect to the rest of the euro area.

Our identification strategy is similar to that of Peersman (2011), who identifies symmetric and asymmetric shocks for the U.K. and euro area by directly imposing cross-country sign restrictions. However, unlike Peersman (2011), who defines asymmetric shocks as those affecting different economies with the opposing sign, we follow a different strategy and argue that asymmetric shocks are more accurately represented by two distinct types of shocks. The first type refers to local (or country-level) shocks, which cannot possibly affect other regions, and the second consists of a narrower subset of common shocks that affect a country and the rest of the euro

area asymmetrically. Our definition aligns with that in an early report by the European Commission: “One Market, One Money” (European Commission 1990). In order to practically implement the identification and estimation of these shocks within BVAR framework, two modifications of a standard model are required. First, to identify country-specific shocks, we impose block exogeneity (or small economy) restrictions, as explained in Comunale and Kunovac (2017) and Deskar-Škrbić, Kotarac, and Kunovac (2020). Second, to identify asymmetric common shocks, we compare the reaction of output in a member country and the rest of the euro area from the estimated historical shock decomposition. Our method, thus, shares similarities with *the narrative sign restrictions approach* described in Antolin-Diaz and Rubio-Ramírez (2018).

We also contribute to the existing literature on formal OCA indices that largely rests on the work pioneered by Bayoumi and Eichengreen (1997, 1998). This framework was extensively used to evaluate the costs of adopting the euro ahead of its launch, and then again later, in the context of candidate countries, as mentioned in Horváth and Komárek (2003), Vieira and Vieira (2012), Skorepa (2013), and Frydrych and Burian (2017). The methodology proposed by Bayoumi and Eichengreen (1997) simply relates nominal exchange rate volatility to some salient OCA features, typically the similarity of business cycles or the strength of trade linkages. Their OCA index is then constructed based on the observation that countries where the nominal exchange rate volatility implied by certain proxies for the OCA criteria is sufficiently low are, supposedly, more prone to abandon their autonomous monetary and exchange rate policy. Clearly, the OCA criteria and the notion of importance of symmetric shocks are not less important for countries *already in a currency union*. Indices developed by Bayoumi and Eichengreen, however, are based on the volatility of nominal exchange rate between a candidate and the currency union and, therefore, are not suitable to keep track of OCA properties of an existing currency union. The absence of formal OCA indices for the euro area is an important measurement gap that we aim to address by constructing a new OCA index, now based on the relative importance of symmetric shocks.

Finally, we contribute to literature on OCA endogeneity for the euro area, a concept advocated by Frankel and Rose (1997), who

challenge the static formalization of traditional OCA theory. There has been no consensus on how euro adoption affects the synchronization of a member with the euro area so far (with limited data availability playing a role in this regard). Interestingly, Campos, Fidrmuc, and Korhonen (2019) conduct a large meta-analysis that is supportive of the OCA endogeneity hypothesis. Willett, Permpoon, and Wihlborg (2010) or Caporale, DeSantis, and Girardi (2015), by contrast, suggest that the conclusion that one needs not worry *ex ante* about optimum currency area conditions, as suggested by the OCA endogeneity theory supporters, is overly benign. To select among these opposing views, a test for the OCA endogeneity hypothesis is required. For this, the literature usually relies on simple correlations between business cycles or, alternatively, on simple ordinary least squares (OLS) specifications used to identify the effect of bilateral trade patterns on income correlations. These approaches suffer from lack of identification and a structural interpretation of the results. Instead, we propose a formal test for OCA endogeneity by studying how the incidence of symmetric shocks evolves over time after a country joins a currency union. In doing so, we use D-I-D methods from Callaway and Sant'Anna (2021) and compare how the incidence of common shocks evolves for new members, in comparison to similar non-euro-area EU European countries. By including a proper control group of countries into the analysis, we address concerns raised in the literature that increased correlation between members may not be a consequence of joining the euro area only (see Willett, Permpoon, and Wihlborg 2010).⁵

3. Modeling Framework

We first specify a structural BVAR model for each euro-area member under analysis to separate local (idiosyncratic) from common shocks underlying GDP growth. We then show how to use this

⁵Willett, Permpoon, and Wihlborg (2010) report: “While the correlations of growth rates among euro area countries rose substantially after the creation of the euro, the correlations of the non-euro area European countries with the euro area countries rose even more, so clearly something besides just the creation of the euro was going on.”

decomposition to separate the impact of symmetric from that of asymmetric shocks on GDP growth.

3.1 A Small Open-Economy BVAR for a Euro-Area Country

We start by specifying an open-economy Bayesian vector autoregression (BVAR) model for 15 euro-area countries⁶ each in turn. The model is thus defined from the perspective of the individual country and includes six variables: GDP growth and the inflation rate of euro-area country i , which is considered the *home* country, the rest of the euro area (REA, the euro area excluding country i), and the rest of the world (RoW).⁷ All details of the data used in this paper are described in Table A.1 of the appendix.

Consider a generic structural VAR with k lags that is represented as follows:

$$A_0 y_t = \mu + A_1 y_{t-1} + \dots + A_k y_{t-k} + \varepsilon_t, \quad t = 1, \dots, T, \quad (1)$$

where y_t is an $n \times 1$ vector of observed variables, A_j are fixed $n \times n$ coefficient matrices with invertible A_0 , μ is an $n \times 1$ fixed vector, and ε_t are structural economic shocks with a zero mean and covariance matrix I_n . The reduced-form VAR model is obtained from (1) by pre-multiplying the equation by $(A_0)^{-1}$:

$$y_t = c + B_1 y_{t-1} + \dots + B_k y_{t-k} + u_t, \quad t = 1, \dots, T, \quad (2)$$

where $B_j = A_0^{-1} A_j$, $c = A_0^{-1} \mu$, $u_t = A_0^{-1} \varepsilon_t$, and $E(u_t u_t') = \Omega = (A_0' A_0)^{-1}$. To identify the structural model in Equation (1) additional restrictions are required. It is important to note that for any $n \times n$ orthogonal matrix Q (i.e., $Q Q^T = Q^T Q = I$), pre-multiplying (1) by Q results in an observationally equivalent structural model, i.e., whose reduced-form representation is also Equation (2). Identification methods relying on *sign restrictions* are generally based on

⁶Our country coverage includes 15 euro-area countries, excluding Malta, Cyprus, Luxembourg, Ireland (due to exceptionally volatile GDP series), and Croatia (which joined the union in 2023).

⁷GDP for the rest of the world is represented by the sum of GDP of Norway, Switzerland, Turkey, Russia, USA, Canada, Mexico, Brazil, Australia, New Zealand, Japan, China, Hong Kong, Korea, and the EU but non-euro-area countries. Inflation is calculated from trade-weighted CPI, from 11 main trading partners of the euro area.

this principle (see, e.g., Canova and De Nicolo 2002, Rubio-Ramírez, Waggoner, and Zha 2010, and Arias, Rubio-Ramírez, and Waggoner 2014, 2018).

To separately identify country-specific shocks from common shocks, it is necessary to impose two types of restrictions. First, to ensure that country-specific shocks cannot affect foreign (REA and RoW) variables, restrictions on the impulse response function at $t = 0$ have to be imposed. In practice, zero restrictions on matrix A_0^{-1} , in addition to sign restrictions, can be implemented as suggested by Arias, Rubio-Ramírez, and Waggoner (2014, 2018). However, these restrictions alone do not prevent country-specific shocks from affecting foreign variables at horizons beyond $t = 0$. To achieve these *small country* or *block exogeneity* restrictions, zero restrictions on additional regression parameters are required. In other words, while in the reduced-form model (2) domestic GDP depends on its own lags as well as lags of foreign variables, foreign variables depend on their own lags only.⁸ To impose zero restrictions on the regression parameters, we rewrite (2) in a more compact form:

$$y_t = X_t' \beta + u_t, \quad (3)$$

where $X_t' = I_n \otimes [1, y_{t-1}', \dots, y_{t-k}']$ and $\beta = \text{vec}([c, B_1 \dots B_k]')$. Within the Bayesian estimation framework, β can be restricted by setting an appropriate prior distribution. The usual choice of the natural conjugate (normal-inverse-Wishart) prior, although beneficial in terms of tractability of the posterior probability density function and computational speed, is not suitable for this purpose because it assumes a Kronecker-type structure of the prior covariance of VAR coefficients that is rather inflexible. In this case, the prior covariances are proportional to each other across equations so that it is not possible to independently set the prior for a subset of parameters in selected equations as would be needed to impose block exogeneity. Kadiyala and Karlsson (1997), Sims and Zha (1998), Koop and Korobilis (2010), or Carriero, Clark, and Marcellino (2019) explain

⁸Kadiyala and Karlsson (1997) and Miranda-Agrippino and Ricco (2018) provide another example where asymmetric treatment of the endogenous variables in a VAR is appropriate and in line with *prior beliefs*: economic theory suggests that *money neutrality* implies that the money supply does not Granger-cause real output.

why standard conjugate priors may appear as an overly restrictive choice for some applications of BVAR models.⁹ Instead, an *independent* normal-inverse-Wishart prior is more appropriate in our case, as the prior beliefs for the VAR coefficients and error covariance matrix are set independently, as also implemented in Deskar-Škrbić, Kotarac, and Kunovac (2020):

$$\beta \sim N(\underline{\beta}, \underline{V}_\beta), \quad \Omega \sim IW(\underline{M}, \underline{\gamma}).$$

Conditional posterior distributions $p(\beta|y, \Omega)$ and $p(\Omega|y, \beta)$ for this prior have the following form:

$$\beta|y, \Omega \sim N(\bar{\beta}, \bar{V}_\beta), \quad \Omega|y, \beta \sim IW(\bar{M}, \bar{\gamma}),$$

where

$$\bar{V}_\beta = \left(\underline{V}_\beta^{-1} + \sum_{t=1}^T X_t \Omega^{-1} X_t' \right)^{-1},$$

$$\bar{\beta} = \bar{V}_\beta \left(\underline{V}_\beta^{-1} \underline{\beta} + \sum_{t=1}^T X_t \Omega^{-1} y_t \right),$$

and

$$\bar{\gamma} = T + \underline{\gamma}, \quad \bar{M} = \underline{M} + \sum_{t=1}^T (y_t - X_t' \beta) (y_t - X_t' \beta)'$$

While the posterior distribution is no longer available in the closed form, conditional posterior distributions are readily available, and a Gibbs sampler is therefore used to draw an approximate sample from the posterior of the reduced-form parameters, β and residual covariance matrix Ω . To evaluate the properties of a simulated sample from the posterior, in Appendix Section A.4 we provide a detailed Markov chain Monte Carlo (MCMC) convergence diagnostics; see Geweke (1992) or Chib (2001).

⁹There are, however, alternative approaches proposed in the recent literature to circumvent this forced symmetry imposed by standard normal-inverse-Wishart prior. For instance, Chan (2019) proposes *asymmetric conjugate priors* that do not preserve the VAR Kronecker structure when forecasting with large Bayesian VARs. An alternative strategy for implementing asymmetric priors is proposed by Carriero, Clark, and Marcellino (2019).

To impose zero restrictions on some regression parameters using an independent normal-inverse-Wishart prior, we assume zero-mean priors with very small variance for all small country parameters in the REA and the RoW equations. For example, to restrict the j -th element of β , we can set $(\underline{\beta})_j = 0$ and $(\underline{V}_\beta)_{jj} = \varepsilon$, with ε being some small positive number. This attaches a large weight to the zero-mean prior parameters when calculating posteriors. Thus, sample information is largely ignored, as the posteriors of these coefficients will be predominantly influenced by the prior. Other elements of $\underline{\beta}$ and \underline{V}_β are set to shrink posterior parameters in the spirit of the Minnesota prior. Hyperparameters are set to $\lambda_1 = 1$, $\lambda_2 = 1$, $\lambda_3 = 1$, and $\lambda_4 = 10^4$, which reflects our choice to use non-informative priors.

The structural BVAR models used in our analysis are all specified in log differences and estimated at a quarterly frequency on period 1996:Q1–2023:Q3 using two lags. Experimenting with different lag numbers did not change our results significantly.

3.2 Shock Identification: Local vs. Common Shocks

We identify six shocks in each of the country BVARs using the sign and zero restrictions as summarized in Table 1. We impose restrictions directly on the impulse response functions on impact only. Additional zero restrictions on autoregressive parameters are set as explained in Section 3.1.

Country-specific aggregate demand and supply shocks affect domestic GDP growth and inflation, but cannot affect real activity and prices in the REA or RoW. A demand shock is associated with positive correlation between GDP growth and inflation, and supply shocks are associated with a negative correlation between the two, as usually assumed in related literature; see, for example, Comunale and Kunovac (2017), Forbes, Hjortsoe, and Nenova (2018), or Bobeica and Jarociński (2019). In addition, only supply shocks can have a long-run impact on GDP, whereas the cumulative response to demand-side shocks is restricted to zero in the long run; see, for example, Blanchard and Quah (1989) and Forbes, Hjortsoe, and Nenova (2018).¹⁰ Finally, by appropriately restricting the

¹⁰A standard view, as implemented in our specification, that all shocks with a permanent effect on output are supply shocks and those with a transitory effect

Table 1. Short-Run and Long-Run Restrictions

Shock/Variable	GDP_{Home}	$HICP_{Home}$	GDP_{REA}	$HICP_{REA}$	GDP_{RoW}	π_{RoW}
	<i>Short-Run Restrictions at $t = 0$</i>					
Local AD (Country Specific)	+	+	0	0	0	0
Local AS (Country Specific)	+	-	0	0	0	0
Common AD (Euro Area)	+	+	+	+	0	0
Common AS (Euro Area)	+	-	+	-	0	0
Common AD (Global)	+	+	+	+	+	+
Common AS (Global)	+	-	+	-	+	-
<i>Long-Run Restrictions</i>						
Local AD (Country Specific)	0	?	?	?	?	?
Local AS (Country Specific)	?	?	?	?	?	?
Common AD (Euro Area)	0	?	?	?	?	?
Common AS (Euro Area)	?	?	?	?	?	?
Common AD (Global)	0	?	?	?	?	?
Common AS (Global)	?	?	?	?	?	?

Note: AD denotes aggregate demand and AS denotes aggregate supply. (+) = positive response; (-) = negative response; (0) = no response; (?) = unrestricted response. GDP_{Home} denotes GDP growth of a euro-area country, GDP_{REA} for the rest of the euro area, and GDP_{RoW} for the rest of the world; $HICP_{Home}$ denotes HICP inflation for a euro-area country, $HICP_{REA}$ for the rest of the euro area, and π_{RoW} inflation for the rest of the world. Details about all variable definitions are listed in Table A.1 of the appendix.

autoregressive coefficients, we also impose that REA and RoW variables must not depend on lagged values of home-country variables. This assumption, together with restrictions imposed on the impulse response function (IRF) at $t = 0$, is sufficient to fully separate local shocks from other shocks at all horizons.

Common euro-area aggregate demand and supply shocks affect macroeconomic indicators (GDP growth and inflation) in the home country and the rest of the euro area, but cannot affect RoW. Initially at $t = 0$, that impact is symmetric, but it may become asymmetric afterward, so later we distinguish between common symmetric and common asymmetric shocks. Demand-side and supply-side common shocks are, respectively, characterized by a positive and negative correlation between GDP growth and inflation. Besides that, only supply shocks can affect GDP in the home country in the long run. We assume that the two euro-area-specific common shocks cannot affect GDP growth or inflation in the rest of the world. To achieve that, we need to impose both restrictions at $t = 0$, but also on corresponding regression coefficients in equations for RoW variables.

Common global shocks are two global shocks that simultaneously affect all regions—the country under consideration, the rest of the euro area, and the rest of the world. Expansionary common global demand shocks initially affect all six variables under analysis positively. Expansionary global supply shocks affect GDP growth positively in all regions and affect negatively consumer inflation across all blocs. We also assume that global common demand shocks cannot affect home GDP in the long run.¹¹

are demand shocks is questioned by recent literature. Furlanetto et al. (2021) attempt to identify demand shocks that can have a permanent effect on output through hysteresis effects; Coibion, Gorodnichenko, and Ulate (2017) point to cyclical sensitivity of mainstream estimates of potential output; and González-Torres, Gumiel, and Szörfi (2023) show how output gap estimates are affected when we allow for supply shocks with transitory effect only. For that reason, different approaches to separate demand versus supply should not be interpreted unconditionally, independently from the underlying identifying assumptions. Separation between symmetric versus asymmetric shocks, based on exact exclusion and sign restrictions—a focus of our paper—is, in contrast, less arbitrary.

¹¹To impose restrictions that only supply shocks may affect GDP in the long run, we impose long-run restrictions on how demand shocks affect domestic GDP only. Ideally, we may want to impose some more restrictions—perhaps also on

3.2.1 *Relative Importance of Shocks: Historical Shock Decomposition*

The relative importance of the identified country-specific and common shocks in individual euro-area countries can be gauged from the historical shock decomposition of the estimated country BVARs. The *relative importance* of shock k to variable j at period t can be calculated from

$$\widetilde{y}_{jt}^k = \frac{|y_{jt}^k|}{\sum_{l=1}^n |y_{jt}^l|}, \quad (4)$$

where y_{jt}^k represents the contribution of shock k to the historical shock decomposition of variable j at period t and n denotes the total number of shocks. Relative importance of individual shocks could alternatively be evaluated using the variance decomposition. However, historical decomposition, by construction, seems to be better suited for tracking the importance of various shocks *over time*.

3.3 *Mapping Local vs. Common to Symmetric vs. Asymmetric Shocks*

In this section we reorganize the identified common and local shocks into symmetric versus asymmetric shocks, which is the decomposition of central interest for our analysis. *Asymmetric shocks* are not only all country-specific shocks but also those common shocks that affect a country and the rest of the euro-area asymmetrically, despite being initially *common* by definition. In fact, some shocks common to the entire euro area may have asymmetric impact on different euro-area members, possibly due to the differences in the initial cyclical states, economic structures, economic behavior, or preferences across the countries; see, e.g., European Commission (1990). Such a definition of (a)symmetric shocks spells out the idea that whenever a country's economy is predominantly driven by country-specific or

how these common demand shocks affect REA and RoW growth, but this is prevented by *the order condition*; see Arias, Rubio-Ramírez, and Waggoner (2014) for an explanation. Nevertheless, we experimented with different combinations of long-run restrictions—for example, instead of always restricting reaction of home GDP, we would restrict GDP of the RoW and REA—which resulted in very similar results.

common shocks with an asymmetric impact, the membership to the monetary union is more costly.

The relative importance of country-specific shocks can be calculated directly using Equation (4). To separately identify asymmetric from symmetric common shocks, we compare the effect of common shocks on the individual country under consideration and REA using historical shock decomposition of GDP in each period. Formally, let *home* and *REA* index any euro-area country and the REA, respectively, and let *k* denote an identified common shock (euro area or global). Then, a common shock *k* is said to be *asymmetric*¹² in period *t* whenever the home country and REA react to this shock with the opposite sign, i.e., $y_{home,t}^k y_{REA,t}^k < 0$. *Symmetric shocks*, in contrast, are those common shocks *k* for which $y_{home,t}^k y_{REA,t}^k > 0$ in period *t*.

Regarding the interpretation of our shock decomposition, the identified common shocks are not the same for all members, as real activity in the REA is different for different home countries. Therefore, the *symmetry*, as defined in our model, is to be seen from the individual country's point of view. Relative importance of symmetric shocks for each country is a country-specific measure of net benefits or *welfare* from sharing the common monetary policy. For example, a shock that affects a country and the rest of the euro area with the same sign is, according to our definition, symmetric for that country, irrespective of a possible mixed response among other currency union members. Those countries that react, at the same time, very differently from the rest of the euro area consequently derive lower

¹²Our definition of asymmetric common shocks is based on *the phase synchronicity* between business cycles measured by GDP growth. This choice, however, is not unique; there are alternative definitions. For example, we may be interested in those asymmetric shocks that affect various countries with the same sign but with very different amplitudes; see, for example, Mink, Jacobs, and de Haan (2012) for a concept of *similarity* between cycles. The COVID-19 shock, for instance, while affecting all the countries negatively, has increased the heterogeneity among euro-area members in terms of the severity of recession (Muggenthaler, Schroth, and Sun 2021). Alternative variants of asymmetric shocks would be more arbitrary compared to ours. In a case in which asymmetric shocks are identified using the similarity of business cycles, a threshold on the level of contribution of shocks should be imposed (or estimated) in order to distinguish between symmetric and asymmetric shocks with our *quasi-narrative* approach. It is, however, not clear how to identify this threshold.

levels of *welfare* upon such a shock occurring, because they are hit asymmetrically and a policy stance of one-size-fits all policy is most likely unsuitable. In that sense, our definition of symmetric shocks departs from a definition where a symmetric shock is assumed to affect *all* the members of a currency union in the same way.

4. Empirical Results

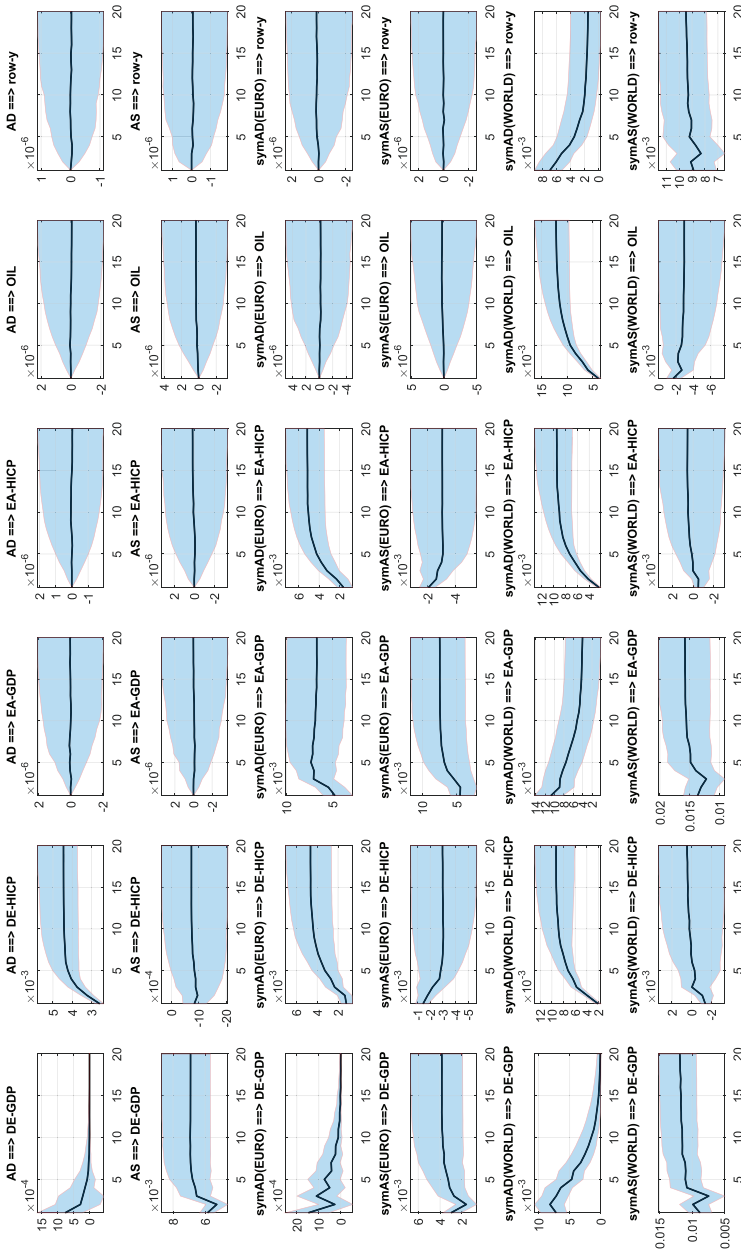
We first review the results from country BVARs in order to show that our identifying assumptions are sufficient to properly distinguish local from common euro-area and global shocks. After that, we study the difference between common and symmetric shocks, as they need not be the same. We finally conduct a detailed *shock accounting* exercise to study drivers of GDP growth in euro-area members. Based on this shock decomposition, we provide a new test for OCA endogeneity and construct a timely OCA index for the euro area.

4.1 BVAR Results—Local vs. Common Shocks

We show impulse responses for Germany (Figure 1) and for other countries in Appendix Section A.2 (Figures A.2–A.5), to verify that the sign and zero restrictions from Table 1, together with additional zero restrictions on autoregressive parameters, are sufficient to separate the local and the two common shocks. Two local shocks in the top rows cannot influence the rest of the euro area or the rest of the world over any time horizon. Likewise, common euro-area shocks in the third and the fourth row cannot affect the RoW. In contrast, the common global shocks in the two bottom rows affect all three *regions* similarly.


Figure 2 compares the relative importance of local, euro-area common, and global common shocks for GDP growth across countries. To summarize the results obtained from identified country BVARs, we also show the cross-country average importance of each shock—both unweighted and (GDP-)weighted versions. At this point, we do not explicitly label each country in the graph and focus on the properties of cross-country distribution. The global common shocks are clearly the most important type of shocks, accounting for the biggest share of all shocks, confirming in particular the small

Figure 1. Impulse Response Functions for Germany
(median and 68 percent posterior bands)



Note: Impulse responses of each variable are in separate columns: German GDP and inflation (GDP and HICP), rest of the euro area GDP and inflation (GDP_{REA} and $HICP_{REA}$), and the rest of the world GDP and inflation (GDP_{ROW} and $INFL_{ROW}$). There are overall six shocks in the rows: local aggregate demand and supply (AD and AS), euro-area-specific symmetric AD and AS ($symAD(EURO)$ and $symAS(EURO)$), and global symmetric AD and AS ($symAD(WORLD)$ and $symAS(WORLD)$).

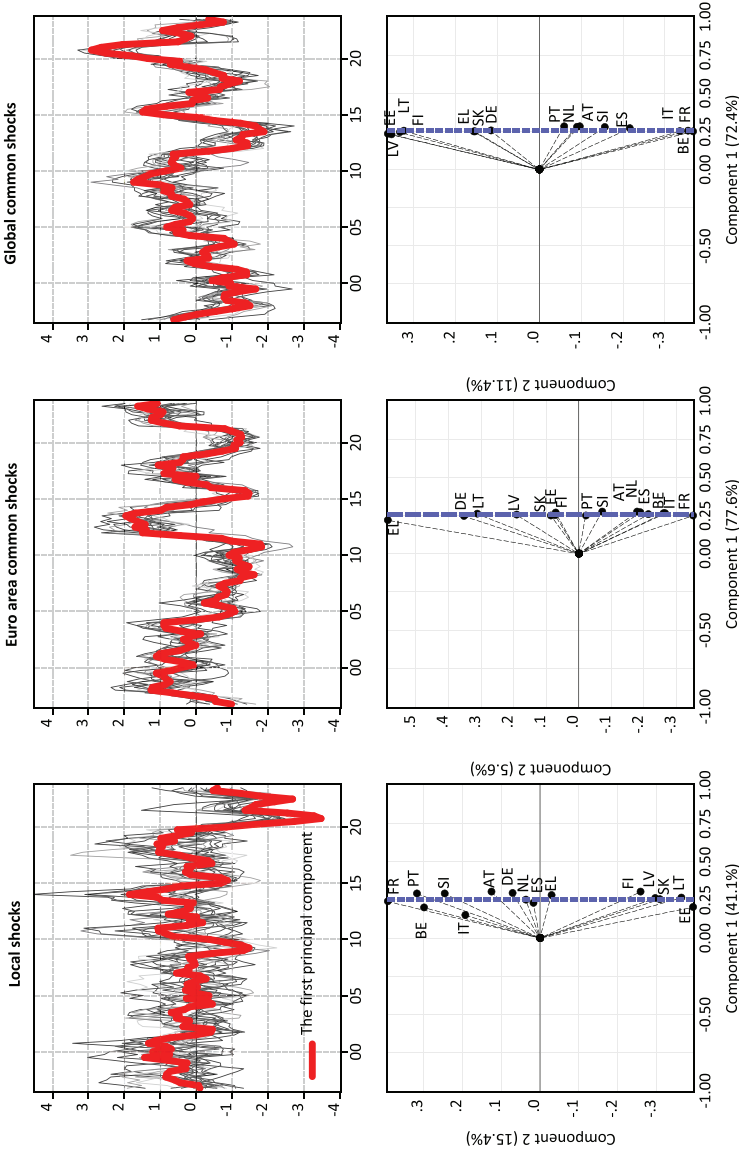
Figure 2. Relative Importance of Local vs. Euro-Area Common vs. Global Common Shocks across Member States



open-economy nature of the euro area. The cross-country variation is overall the largest for local and global shocks, but diminishing over time. For local shocks, the unweighted average hovers around 30 percent, above the weighted average. Conversely, for global shocks, the unweighted average lies below the weighted average. For both local and global common shocks, the difference between the two averages has been closing in over the sample, mostly driven by smaller, *new* euro-area members that are being increasingly more exposed to global symmetric shocks and less so to idiosyncratic economic disturbances. Also the dispersion of global shocks across the countries reduced significantly in the last decade, but the share of global shocks was much less stable, and peaked during the pandemic shock. As regards the euro-area-specific common shocks, they are less dispersed and generally stable in importance over the observed period. In line with the usual narrative, they peak during the euro debt crisis. Non-negligible importance of common euro-area shocks—also illustrated by yellow contributions to GDP growth in Figures A.6–A.9—supports the existence of a European business cycle emerging from our analysis. This has been discussed intensively in the literature, but without reaching a firm consensus.

Results in Figure 2 are informative about *the level* of the relative importance of various shocks over time and across countries. However, in order to better understand the comovement behind estimated importance of shocks—a feature that may be overlooked easily when looking at Figure 2 only—we rely on the principal component analysis (PCA), based on standardized data; see Figure 3 (upper panel). After the transformation, the high dispersion of local shocks across countries still persists, while the relative importance of global shocks appears highly correlated across countries. Euro-area-specific contributions were already very little dispersed before standardization, and this transformation had little impact here. A simple PCA, applied within our methodological framework, can illustrate the commonality among euro-area members very well. The first principal component explains 77.6 percent and 72.4 percent of the common variation for euro-area and global common shocks, respectively, in contrast to only 41 percent for local shocks (x-axis, Figure 3, lower panel). To illustrate that all countries load similarly to first components of common shocks, we perform a simple visual test and compare calculated country loadings (the black dots) with

Figure 3. Relative Importance of Local vs. Common Shocks:
A Principal Component Analysis



Note: The relative importance of shocks in the upper panel is standardized to have zero mean and unit standard deviation. Blue vertical lines in the lower panel refer to the average loadings of the first principal component. The black dots refer to the loadings of individual countries to the first principal component.

the cross-country average loading (the blue vertical lines). Our simple test suggests that no country deviates systematically from the common cycles; only Greece exhibits a slightly larger deviation from the European cycle.

There is some remaining commonality captured by the second principal component, 5.6 percent for euro-area common and 11.4 percent for global common shocks, shown on the y-axis. Interestingly, some country groups emerge to form clusters. In particular, there is evidence of clustering patterns according to *geographical proximity*—the Baltic countries and Finland form a cluster in the case of global common shocks. These countries also show up among the countries that are less affected by common euro-area shocks during the debt crisis—Greece (EL), Germany (DE), Finland (FI), and the Baltics (EE, LT, LV). This is not to be confused with a wrong conclusion that Greece was not affected by the sovereign debt crisis. It was strongly affected, however mostly through local idiosyncratic shocks that would later spill over to other members, but not through euro-area common shocks, as was the case for other countries, such as Italy, Portugal, or Spain. Our empirical model is capable of properly differentiating these types of shocks. Local shocks are unrelated across countries by construction. Nevertheless, some clustering evidence emerges that relative importance is less important for old euro-area members compared to new members in the case of local shock loadings.

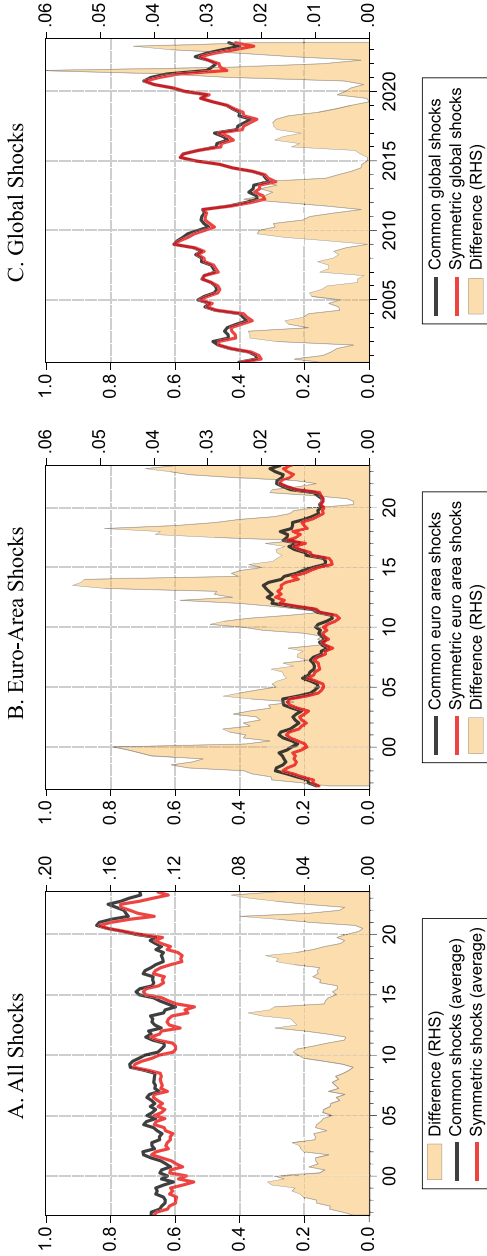
4.2 *Shock Accounting: Symmetric vs. Asymmetric Shocks and Implications for Monetary Policy*

After we explain the difference between common and symmetric shocks, we show the results of our *shock accounting* exercise. First we document some differences in composition of shocks between countries and then, finally, study relative importance of individual shocks for the euro area as a whole and implications for monetary policy.

4.2.1 *Common and Symmetric Shocks: A Comparison*

Asymmetric shocks, as explained in Section 3.3, include both local shocks and asymmetric common shocks. Equivalently, not all common shocks are symmetric. Figure 4 compares cross-country average

Figure 4. Common and (a)symmetric Shocks



Note: Cross-country average share of symmetric shocks is represented by red lines; cross-country average share of common shocks is represented by black lines. The difference between the two measures is represented by the yellow area, measured against the right-hand axis.

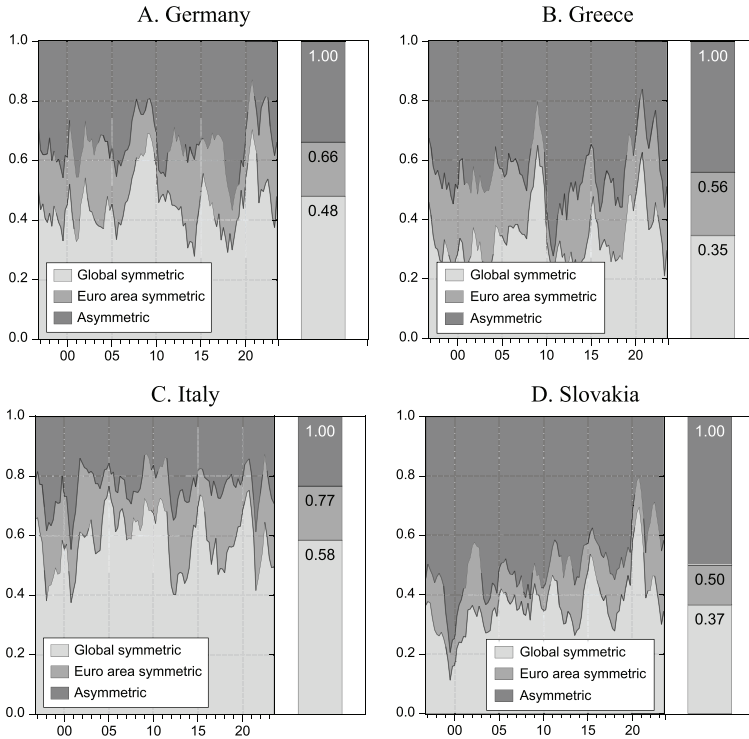
of common and symmetric shocks using our definitions and suggests that the difference between the two is small. Some common shocks could indeed be asymmetric, but never more than 8 percentage points (pp) of overall variation in GDP growth, with only 3 pp on average over the observed period; see Figure 4A. Most of the asymmetric common shocks are identified as euro-area common shocks (mean = 2.2 pp, max = 5.5 pp), less often global (mean = 1.1 pp, but higher max = 6.2 just after COVID); see Figure 4, panels B and C. European asymmetric common shocks may intensify just after the periods of increased commonality: notably, they peak after the euro debt crisis. Global asymmetric shocks, similarly, peaked right after the pandemic, suggesting that major crises tend to evoke uneven reactions across countries.

4.2.2 *Shock Decomposition by Country*

The relative importance of symmetric versus asymmetric shocks is not always homogeneous across euro-area members. Typically, factors that drive the differences among them are related to how previous crises (GFC or euro debt crisis) affected different members, when they join the euro area (*old* or *new* members), or how individual members align with global developments in general.

To better understand the differences observed across countries, Figure 5 compares the relative importance of symmetric versus asymmetric shocks for four euro-area countries, capturing distinct and particularly interesting shock configurations: Germany, Greece, Italy, and Slovakia. Greece is a member with a generally large share of asymmetric shocks—around 44 percent of overall variation in GDP growth; see Figure 5B. These shocks were especially important around the European sovereign debt crisis, when they dominated and drove 70 percent of the overall business cycle variation. Unlike in other countries, the macroeconomic deterioration was due to local shocks that later spread throughout Europe. However, later on, Greece has been increasingly affected by common symmetric shocks. In Germany, the largest euro-area economy, we also find a fairly large share of asymmetric shocks; see Figure 5A. This could come as a surprise as, by its sheer size, the German economy is well correlated to the overall euro area. It is, however, less well correlated with the rest of the euro area. Indeed, due to its size, a difference in

Figure 5. Shock Contributions in Select Member Countries



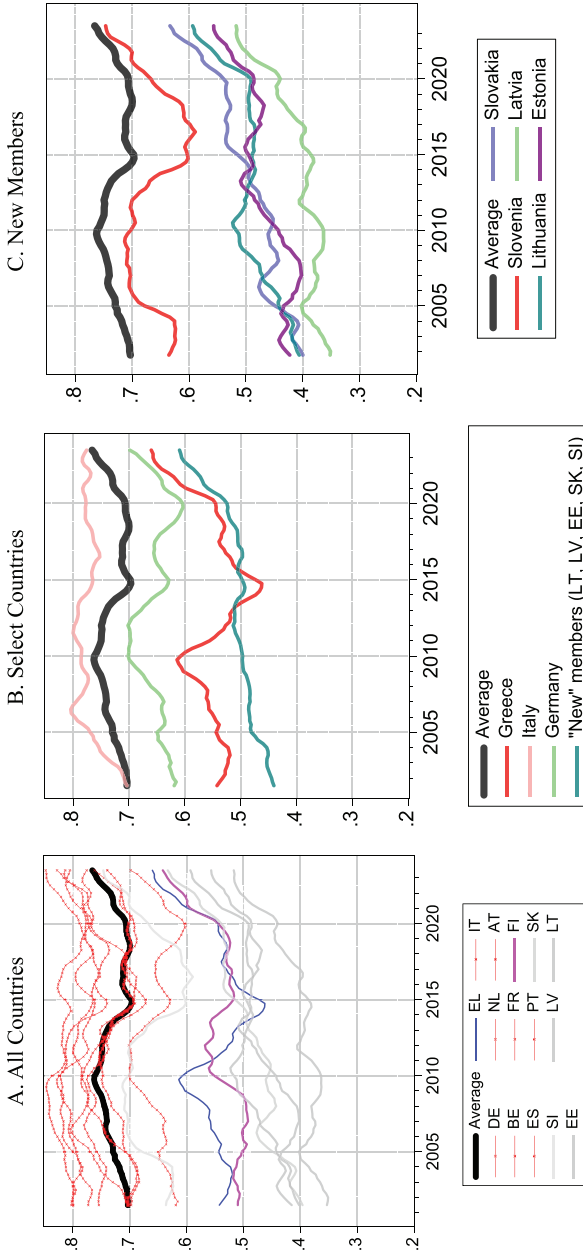
Note: Sample average for each shock is presented in bars on the right (cumulative).

correlation of Germany to total euro area in contrast to correlation to rest of the euro area (i.e., Germany excluded) is sometimes large; see Figure A.1, panel C. Our framework recognizes this distinction and occasionally points to a relatively large share of asymmetric shocks. Asymmetric shocks—in contrast to Greece, for example—have often been due to stronger performance of the German economy compared to the REA. One such episode was the period immediately following the GFC, when German GDP, in contrast to the rest of the euro area, recovered exceptionally fast. A detailed explanation of a fast recovery in Germany can be found, for example, in an International Monetary Fund report (IMF 2011). Historical shock decomposition

of GDP growth in Italy (Figure 5C) suggests that common symmetric shocks dominate throughout the sample period even more. In fact, according to the results from our BVARs, the relative importance of symmetric shocks here mostly exceeds those of all other countries in our sample. This is in line with Belke, Domnick, and Gros (2017) and Ferroni and Klaus (2015) who, somewhat to their own surprise, also find very strong cyclical coherence between Italy and other euro-area countries. Local shocks have gained importance in driving GDP growth only very sporadically. The three exemplary countries are compared to aggregate euro area in Figure 6B. Our last example of Slovakia (Figure 5C) shows the typical behavior of a *new* member that is catching up to old member countries, while starting off with a very large share of local shocks that would steadily diminish. Figure 6C compares the weighted average of all member countries with the relative importance of symmetric shocks for other new members as well. Indeed, a clear pattern emerges that all new members (except for Slovenia, which is already in a more mature stage of the same process) are catching up with the euro-area average. An interesting question we also ask is to what extent is this convergence affected by the euro adoption.

Four typical countries illustrate how profiles of relative importance of symmetric versus asymmetric shocks vary across countries and over time. There is another source of possible dissimilarity of shocks if the type of shock—demand or supply—is taken into consideration. Figure A.15 shows the relative importance of asymmetric shocks (local) versus symmetric shocks (euro area or global), decomposed further into demand- versus supply-side shocks. It can be seen that among symmetric shocks, those of a global nature dominate. Asymmetric shocks are relatively more important than euro-area symmetric shocks, with a number of countries being exposed to local shocks exceeding 50 percent of all shocks. Local and euro-area symmetric shocks tend to be more often supply-side shocks than demand-side shocks. In the case of global shocks, demand- and supply-side shocks are more or less equally frequent. As already seen in a less refined decomposition onto local versus common shocks in Figure 2, dispersion decreases for local shocks, mainly driven by new members. Importantly, the crisis periods display distinct mean and standard deviation profiles of the relative importance of the shocks. Around the European debt crisis the relative importance of

Figure 6. Symmetric Shocks: Some Country Differences



Note: All series are transformed to five-year moving averages. In panel A, all *new* members are represented by the gray lines and all other countries, except Finland and Greece, are in red.

symmetric euro-area supply shocks increased significantly as well as the cross-country heterogeneity. The GFC and pandemic were global crises, where the importance of symmetric shocks rose but with very low dispersion.


4.2.3 *Shock Decomposition for the Euro Area and Implications for Monetary Policy*

While relative importance of symmetric versus asymmetric shocks at the country level is important, the focus of *one-size-fits-all* monetary policy is on euro area as a whole. For that reason, we aggregate results from country BVARs and calculate cross-country average relative importance of each type of shock. Results from individual countries are aggregated using GDP as weights.

Figure 7A compares the relative importance of each shock for the euro area as a whole and points to some interesting results. First, symmetric shocks—global in blue and euro-area-specific in red—are the dominant drivers of GDP growth for the euro area as a whole. On average, across countries and over time, they explain a large portion of overall variability in GDP growth—around 73 percent. Among the symmetric shocks, most of them, around 52 pp, are global shocks, and another important part of them, around 19 pp, is due to euro-area-specific shocks; see decomposition in Figure 7B. This finding on the dominance of global symmetric shocks explains why monetary policy of the ECB is often well synchronized with that of other major central banks. Also, the finding that a sizable portion of the euro-area activity is driven by common euro-area shocks, together with insights from historical shock decomposition at a country level (Figures A.6–A.9), provide additional evidence of the existence of the euro-area-specific cycle.

The predominance of symmetric shocks per se is not sufficient for monetary policy to stabilize the union—the nature of the symmetric shocks matters as well. Symmetric shocks are, therefore, regrouped in Figure 7C in order to compare relative importance of symmetric demand versus supply shocks. Both demand (with 32 percent) and supply (with 41 percent) are, on average, important drivers of overall variability of business cycle of the euro area. A large portion of the comovement explained by supply shocks, however, points to possible common inflationary pressures that monetary policy may not be in

Figure 7. Shock Accounting: Symmetric vs. Asymmetric Shocks



a good position to battle against. As a consequence, an overall large share of symmetric shocks does not necessarily signal an increased potential for common policy to stabilize a currency union.

From a country perspective, of either a candidate country or an existing member, relinquishing its own monetary policy entails higher costs in an environment of more prevalent idiosyncratic demand shocks. Presumably, monetary policy is best suited to deal with these types of shocks and, therefore, their dominance is related to larger costs of not having autonomous monetary policy. The decomposition presented in Figure 7A is suitable to address this concern. Importantly, we show that these shocks are a minor driver of business cycles across members of the union—only around 5 percent of overall variability in GDP growth is due to domestic demand. These findings support a view that the costs of abandoning autonomous country-level monetary policy are likely to be contained. Indeed, the results indicate that, on its own, the country's monetary policy would largely face the same type of shocks as countries within the union or, occasionally, when facing local shocks, these would often be local supply shocks, for which the costs of absence of own monetary policy are considerably lower.

Finally, as a by-product, our framework can also be used to gauge a measure of output gap—an important input for monetary policy to understand the link between real economy and inflation. One approach to estimate output gap is accumulating the contribution of demand shocks from BVARs; see Camba-Mendez and Rodriguez-Palenzuela (2003), Coibion, Gorodnichenko, and Ulate (2017), or Chen and Gornicka (2020), for example. However, this is based on a standard, but fairly strong, assumption that all supply shocks have permanent effects on output; see Blanchard (2018) for explanation. In our application, output gap can then additionally be decomposed in terms of local, euro-area, or global shock contributions. We show in Figure A.14 that the euro-area output gap is dominantly being driven by symmetric global shocks according to our model.¹³ Our estimate resembles a standard Hodrick-Prescott approach relatively

¹³A simple variance decomposition exercise suggests that global symmetric, euro-area symmetric, and local shocks explain 77 percent, 17 percent, and 6 percent, respectively, of overall variation in our measure of output gap; see Lindeman, Merenda, and Gold (1980) for methodology used.

well before the GFC and less so afterward, especially during and after the COVID episode, and these differences are likely related to our definition of demand and supply where all demand shocks are cyclical and potential output is being affected by supply only.

4.2.4 Testing for OCA Endogeneity

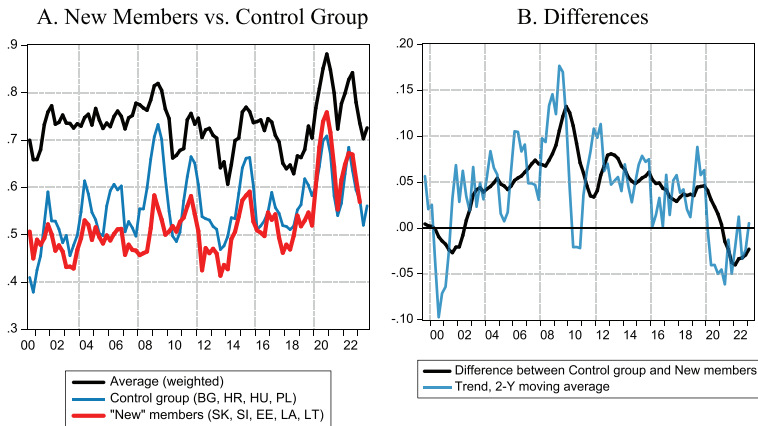
Our results, as shown in Figure 6, illustrate that new member countries, in terms of the incidence of symmetric shocks, converge to the euro-area average. We are now interested in *the euro effect*: how joining the euro area affected the importance of symmetric shocks in these Central and Eastern European countries that all joined the union after 2007—Slovakia, Slovenia, Estonia, Latvia, and Lithuania (SK, SI, EE, LV, LT). For that purpose, by using our identification we first calculate the relative importance of symmetric shocks for these countries and, then, visually compare how the incidence of symmetric shocks has been evolving in comparable non-euro-area countries.¹⁴ After that, to quantify this impact more formally, we use both a standard difference-in-difference regression (denoted D-I-D) as well as an approach from recent advances in the difference-in-difference literature with complex treatment timing; see Callaway and Sant’Anna (2021).¹⁵

A visual comparison of the overall dynamics of symmetric shocks in new members with a *control group* consisting of non-euro EU countries—Bulgaria, Croatia, Hungary, and Poland—resulted in some interesting findings. Figure 8A suggests that both new members and control group are catching up to the euro area over the

¹⁴Our sample ends in 2023:Q3 and Croatia joined the euro area in January 2023. For that reason, we ended our estimation sample a bit earlier, in order to not include post-euro-adoption period.

¹⁵A standard D-I-D estimation based on two-way fixed-effects regression works nicely when dealing with a standard two-periods two-units problem where the untreated group never participates in the treatment, and the treated group becomes treated in the second period. However, there are drawbacks to this approach in a case with more than two time periods and where different units can become treated at different points in time. For example, in our case, a standard D-I-D regression compares the effect of euro adoption with countries that already adopted the euro (and thus maybe behave differently) and those that never adopted the euro (never treated) and those that have not adopted the euro yet. In contrast, the Callaway and Sant’Anna method takes these considerations into account.

Figure 8. Relative Importance of Symmetric Shocks by Country Groups




Note: Relative importance of symmetric shocks for control group is calculated using the methodology explained in Section 3.2.

sample. This is not surprising—all these Central and Eastern European countries went through a similar economic transition and reorientation process toward Western Europe. As a result, the incidence of symmetric shocks has increased in both groups. We are, however, interested in how joining the euro area affected (possibly sped up) this catching-up process. Figure 8B documents the difference between the importance of symmetric shocks in our control group of countries and in new members and provides an informal assessment of the euro effect. Initially, since the euro inception, symmetric shocks in the control group were somewhat more prevailing compared to new members, and the difference between the two groups was trending upward. Over that period, all new members were still outside the euro area. However, it seems that the period when this trend turned broadly coincided with the period when the first new members started joining the euro area.

Figure 9 and Table A.2 provide formal evidence of the euro effect, based on estimated difference-in-difference specifications. In order to understand better these effects, in addition to how all symmetric shocks are affected by euro adoption, we also estimate effects

Figure 9. Difference-in-Difference: Two-Way Fixed Effects (D-I-D) and Callaway and Sant’Anna Method (CS)



Note: The figure compares total effects based on alternative difference-in-difference estimators: standard two-way fixed effects (D-I-D) and the Callaway and Sant’Anna method (CS). D-I-D and CS estimates are shown together with a two-standard-deviation interval. For the CS method, we show estimated heterogeneous effects for different entry periods: 2007 (Slovenia), 2011 (Slovakia), 2011 (Estonia), 2014 (Latvia), and 2015 (Lithuania).

for euro-area-specific and global symmetric shocks separately. For monetary policy purposes, it is also important how predominance of different types of symmetric shocks—demand and supply side—changes after a country adopts the euro. Based on the method by Callaway and Sant’Anna, we find some evidence of *the euro effect*: the incidence of symmetric shocks in new members would on average increase by 4.4 pp as a consequence of joining the euro area. This is almost entirely driven by global symmetric shocks, while the contribution of euro-area-specific shocks is of much smaller magnitude and statistically insignificant; see Figures 9A and B. Standard D-I-D estimates, in contrast, find no significant *euro effect*. Using the Callaway and Sant’Anna method, we also provide evidence of heterogeneity in the euro effect: compared to our control group, not every country benefits equally from joining the euro area in terms of the incidence of symmetric shocks (Figure 9). While some new member countries that first joined the euro area, Slovakia and Slovenia, seem to have benefited from joining the euro area, the findings are more mixed for the Baltics. We found no gains in terms of incidence of symmetric shocks for Estonia and Latvia, and in Lithuania we documented increased incidence of global shocks, but this is largely offset by decrease in euro-area-specific component; see Figure 9B. Finally, a detailed decomposition of estimated euro effect in Table A.2 suggests that increase in dominance of symmetric shocks has been working through both increased incidence of demand and supply, but demand dominates.

Our empirical analysis is concerned with how incidence of symmetric shocks changes after euro adoption in new member countries that joined the union in later stages of enlargement. A similar formal evaluation of the euro inception in January 1999 on shock synchronization, due to a relatively short pre-euro sample, would be somewhat unreliable.¹⁶ There is, however, evidence in related literature that no significant change in business cycle similarity among member states associated with the EMU can be detected (Giannone, Lenza, and Reichlin 2008). A similar, albeit informal, conclusion

¹⁶We have included data starting in 1995—that is the earliest starting date possible from the harmonized Eurostat database—as well as extending the data up to 2023:Q3. Our estimates of the relative importance of shocks start only from 1996:Q4, as we lose some initial data due to taking moving averages.

is suggested by Figure 6A (red lines): symmetric shocks seem to be dominant drivers of GDP growth in founding member countries even before euro inception—they account for almost 80 percent of all shocks. Euro inception itself, as suggested by the literature, probably has little additional impact.

Overall, our analysis suggests that the incidence of symmetric shocks has been increasing for all former socialist European countries, irrespective of their euro-area membership, mostly as a part of a general integration with Western Europe and, more broadly, through ongoing globalization processes. We find, however, that in addition to these common factors that predominantly drive the catching-up process, participating in a currency union, on average, additionally speeds up the process of integration. This greater integration is achieved dominantly through larger incidence of symmetric global demand shocks, suggesting more efficient international trade and investment once adopting the euro, as predicted by early OCA endogeneity literature. This has important implications for candidate countries that consider joining the union: the OCA endogeneity should not be advocated unconditionally and should be necessarily evaluated on a country-by-country basis.

4.3 *Constructing a New Optimum Currency Area Index for the Euro Area*

Cross-country average relative importance of symmetric shocks, as calculated using our framework, is a *shocks-similarity* measure itself, suitable to track OCA properties of the euro area. We argue that dispersion of country-level importance of symmetric shocks also matters. For example, in a case in which the average importance of symmetric shocks is high but heterogeneity across members may have also increased, it may not be sufficient to consider only one of them, first or second moment, in isolation. Indeed, if, for example, high average relative importance of symmetric shocks of 70 percent is attained within a highly heterogeneous environment—the share of symmetric shocks amounts to 90 percent in some countries and only 20 percent in others—common policy will not be optimal for all, with risks of causing political tensions or even possible breakup. We show that a simple *signal-to-noise* ratio of estimated shares of

symmetric shocks takes these considerations into account and qualifies as a reliable OCA index of the euro area.

4.3.1 *Signal-to-Noise Ratio: A Definition*

We construct our OCA measure for the euro area starting from the observation that the common monetary policy will be more successful in stabilizing the euro-area economy if the percentage of symmetric shocks is sufficiently high across member countries. Only in that case will the union-wide monetary policy be well tailored for all member countries. Through the lens of our empirical framework, we additionally argue that the euro area may be *closer to optimal* as a currency area if two conditions are met jointly:

- (i) *The cross-country average of the relative importance of symmetric shocks, denoted μ , is high.* A high value of μ reflects that business cycles across member states are predominantly driven by symmetric shocks, ensuring higher chances for common monetary policy to stabilize all euro-area members simultaneously.
- (ii) *The cross-country standard deviation of the relative importance of symmetric shocks, denoted σ , is low.* In addition, it is desirable that symmetric shocks be of similar importance for all euro-area countries such that the value of σ is low. This implies that for a given value of μ , the summarizing index should penalize high dispersion of the relative importance of symmetric shocks across countries.

The concept of a simple *signal-to-noise* ratio¹⁷ intuitively embeds these two requirements simultaneously:

$$SNR(t) = \frac{\mu(t)}{\sigma(t)},$$

¹⁷In electronics, the ratio of desired electronic signals to unwanted noise, often expressed in decibels (dB) is routinely analyzed to evaluate the signal quality. Here, some analogy can be drawn with our application where high average importance of symmetric shocks cannot provide *clear signal* of favorable OCA properties whenever surrounded by a lot of *noise*.

where $\mu(t)$ and $\sigma(t)$ denote the cross-country sample mean and standard deviation of estimated relative contributions of symmetric shocks, denoted $y_1(t), \dots, y_n(t)$, calculated for a group of n euro-area members:

$$\mu(t) = \frac{1}{n} \sum_{i=1}^n y_i(t) \quad (5)$$

$$\sigma(t) = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (y_i(t) - \mu(t))^2}. \quad (6)$$

In this version of the index, equal weights are attached to each country by definition, so we call it the *unweighted OCA index*. We also compute a *weighted OCA index*:

$$SNR^w(t) = \frac{\mu^w(t)}{\sigma^w(t)},$$

where each country enters the formula for the mean and standard deviation first weighted by its GDP:

$$\mu^w(t) = \sum_{i=1}^n w_i y_i(t) \quad (7)$$

$$\sigma^w(t) = \sqrt{\frac{\sum_{i=1}^n w_i (y_i(t) - \mu^w(t))^2}{1 - \sum w_i^2}}, \quad (8)$$

where $w_i \in [0, 1]$, $\sum w_i = 1$ denote the relative weight of the GDP of country i in the aggregate euro-area GDP. The two measures, $SNR(t)$ and $SNR^w(t)$, have a somewhat different interpretation. When calculating the weighted index $SNR^w(t)$, which accounts for the country sizes, both moments, $\mu^w(t)$ and $\sigma^w(t)$, and the resulting signal-to-noise ratios are by construction dominated by large countries. $SNR^w(t)$ is therefore better at measuring the potential for the common monetary policy to stabilize the overall euro-area economy. On the other hand, economic homogeneity across *all* euro-area members is better reflected in high values for unweighted $\mu(t)$ and low values for $\sigma(t)$.

Signal-to-noise ratios are non-negative numbers and unbounded from above. In our specific case, the relative contributions of symmetric shocks $y_1(t), \dots, y_n(t)$ are all within the interval $[0, 1]$,

implying that the sample means $\mu(t)$ and $\mu^w(t)$ also lie within the same interval. $\sigma(t)$ is bounded by zero from below and, according to Bhatia-Davis inequality (Bhatia and Davis 2000), always bounded from above. For a given average importance of symmetric shocks, irrespective of how $y_1(t), \dots, y_n(t)$ are distributed, Bhatia-Davis inequality states that


$$\sigma(t) \leq \sqrt{\mu(t)(1 - \mu(t))}.$$

Bhatia-Davis inequality provides an upper bound for the standard deviation $\sigma(t)$ and a lower bound for signal-to-noise ratio $SNR(t)$ for a given estimated $\mu(t)$ and, thus, may help to compare estimated values of OCA indices to some known boundary values.

4.3.2 *Signal, Noise, and Signal-to-Noise Ratio*

Figure 10 compares aggregate first ($\mu(t)$ and $\mu^w(t)$) and second ($\sigma(t)$ and $\sigma^w(t)$) moments for the weighted and unweighted samples. The weighted average importance of symmetric shocks $\mu^w(t)$ is always greater than the unweighted figure $\mu(t)$, reflecting the fact that the large euro-area members have, on average, had more coherent business cycles with the rest of the euro area than the other countries. The two mean statistics in Figure 10A peak around global crises: the GFC and the COVID crisis. Importantly, in contrast to very similar short-run dynamics of the two mean series, they have different trending properties: unweighted statistics $\mu(t)$ has had a clear upward trend, largely driven by the catching-up process in new member countries, as explained in Section 4.2.4. In contrast, weighted mean, although already at very high level, has been stagnating, or in the case in which the COVID period is excluded, even slightly decreasing over the observed period. As a consequence, the difference between the two mean measures has been decreasing, as illustrated in Figure 10C. Two dispersion statistics in Figure 10B also show different dynamics over the period under analysis. Unweighted standard deviation $\sigma(t)$ is always larger than weighted statistics $\sigma^w(t)$, suggesting that larger member countries deviate less from cross-country average relative importance. However, dynamics of unweighted statistics is dominated by a strong downward trend, mostly driven by the catching-up process in new (and small) member countries. Weighted dispersion, on the other hand,

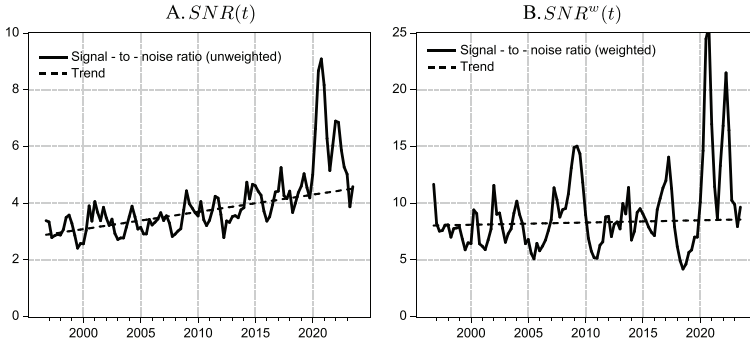
Figure 10. Cross-Country Average and Standard Deviation of Relative Importance of Symmetric Shocks



as a result of a relatively high and stable share of symmetric shocks in large members (see Figure 6A for example) is fairly stable over a longer horizon. In the short term, dispersion of weighted standard deviation is affected differently in different crises: it bottoms out during global crises (the GFC and COVID) and peaks during the European debt crisis. More recently, in the period after the European debt crisis, the gap between weighted and unweighted dispersion has been much smaller compared to the pre-crisis period; see Figure 10C. This trend reflects increased synchronization of small and new members with the rest of the euro area. Previously more idiosyncratic business cycles in these countries were important drivers of the relatively high unweighted dispersion throughout the period before the crisis.

OCA indices $SNR(t)$ and $SNR^w(t)$, shown in Figure 11, are then constructed directly as ratios of $\mu(t)$ and $\sigma(t)$ and $\mu^w(t)$ and $\sigma^w(t)$, respectively. The large values for both the weighted (from 5 to 25) and unweighted versions (from 2 to 9) of the proposed indices provide a clear signal of the average importance of symmetric shocks across countries. The optimality of the euro area as a currency union, as measured by the proposed indices, varies over time and, by the nature of their construction, depends on the types of shocks hitting member countries—local, euro-area-specific, or global. The two indices differ in that the weighted index is always above the unweighted version, mostly reflecting relations between weighted and unweighted moments ($\mu^w(t) > \mu(t)$) and second moments ($\sigma(t) > \sigma^w(t)$). Short-run dynamics, on the other hand, of both signal-to-noise ratios is largely driven by variation of second moments $\sigma(t)$ and $\sigma^w(t)$ that is much larger compared to that of cross-country averages; see Figure A.16. Regarding their long-term dynamics, unweighted measure $SNR(t)$ is trending upward, while weighted measure $SNR^w(t)$, already at a much higher level, remains stable in the long run. Overall, constructed OCA indices suggest that the OCA features of the euro area, despite the documented catching-up process of new members, are stagnating in the longer run at levels very similar to those from the pre-euro era.

Our results clearly illustrate that main crises have different mean and standard deviation profiles; see Figure A.15, for example. The European debt crisis was characterized by increased importance of symmetric euro-area shocks and increased cross-country


Figure 11. Signal-to-Noise Ratio

heterogeneity. COVID and the GFC, on the other hand, were global crises, where the importance of global symmetric shocks was elevated, but the dispersion across countries was very low. A useful measure of monetary policy potential should, therefore, take these observations into account. Our indicator is able to capture different types of crises, as it is the *signal* with respect to surrounding *noise*, $\mu(t)/\sigma(t)$, and not necessarily the signal $\mu(t)$ or $\sigma(t)$ alone, that is crucial for stabilizing potential of a common monetary policy. To illustrate that, Figure 12 compares $\mu(t)$, $\sigma(t)$, and ratio $SNR(t)$ for all symmetric shocks (first column), euro-area-specific shocks (middle column), and global shocks (third column). Two global crises—marked with red vertical lines—are characterized by a large share of common global shocks surrounded by low dispersion. Consequently, signal-to-noise ratio, in the third row, increases strongly and provides a clear signal for policy stance. In contrast, the euro-area debt crisis, marked with blue lines, increased overall incidence of symmetric shocks, but increased volatility. As a result, signal-to-noise ratio remained largely unchanged, despite the increase in symmetric shocks.

5. Conclusions

This paper investigates the importance of asymmetric shocks on the euro area and their implications for monetary policy within the currency union. By employing a structural, sign- and zero-restricted, open-economy BVAR model, we are able to differentiate between

Figure 12. Turning Points in Averages, Standard Deviations of Symmetric Shocks and in Signal-to-Noise Ratios



Note: Red vertical lines refer to global crises; blue vertical lines refers to the European debt crisis.

symmetric and asymmetric shocks and assess their relative importance in shaping euro-area business cycle dynamics. Our findings indicate that, while symmetric shocks are the primary drivers of these cycles, the OCA features of the euro area are stagnating in the longer run at levels very similar to those from the pre-euro era.

We introduce a new OCA index underscoring that for a common monetary policy to be effective, there must be a high but not excessively varied importance of symmetric shocks across countries. This index proves particularly valuable in distinguishing between global and localized crises, supporting stability assessments within the euro area.

Our results are also favorable for the OCA endogeneity hypothesis, showing that euro adoption increases the incidence of symmetric shocks, also fostering greater business cycle synchronization among new member states. This challenges the traditional belief that a high degree of pre-adoption coherence is essential for euro membership, suggesting that convergence can continue post-adoption.

The OCA index offers key insights for policymakers by indicating conditions under which common monetary policy is most effective and identifying challenges posed by asymmetric shocks. At the same time, our findings also highlight the necessity of ongoing structural convergence among member economies to ensure the stability of the union. Overall, our research affirms the euro area's viability as a monetary union, whereby the benefits of membership, particularly through post-adoption synchronization, outweigh the costs. Nevertheless, continuous monitoring and policy adjustments remain essential to address asymmetric shocks and sustain long-term stability.


Appendix

A.1 Data

Table A.1. Variable Descriptions

Variable	Definition	Source
Real GDP (Country)	Chained-link volume, million of euros	Eurostat
Real GDP (Rest of the Euro Area)	Euro-area aggregate excl. one member country	Eurostat
Real GDP (World)	Sum of real GDP of Norway, Switzerland, Turkey, Russia, USA, Canada, Mexico, Brazil, Australia, New Zealand, Japan, China, Hong Kong, Korea, and EU but non-euro-area countries	Eurostat, OECD
HICP (Country)	Harmonized index of consumer prices, index-weighted sum of HIPC indices excl. one member country	Eurostat
HICP (Rest of the Euro Area)		
CPI (World)	Trade-weighted inflation, using data for United States, United Kingdom, China, Switzerland, Poland, Czech Republic, Sweden, Hungary, Japan, Denmark, India	IMF (CPI), Eurostat (Trade Weights)

Figure A.1. GDP Growth: Cross-Country Correlation



A.2 Impulse Response Functions

Figure A.2. Impulse Response Functions for France
(median and 68 percent posterior bands)

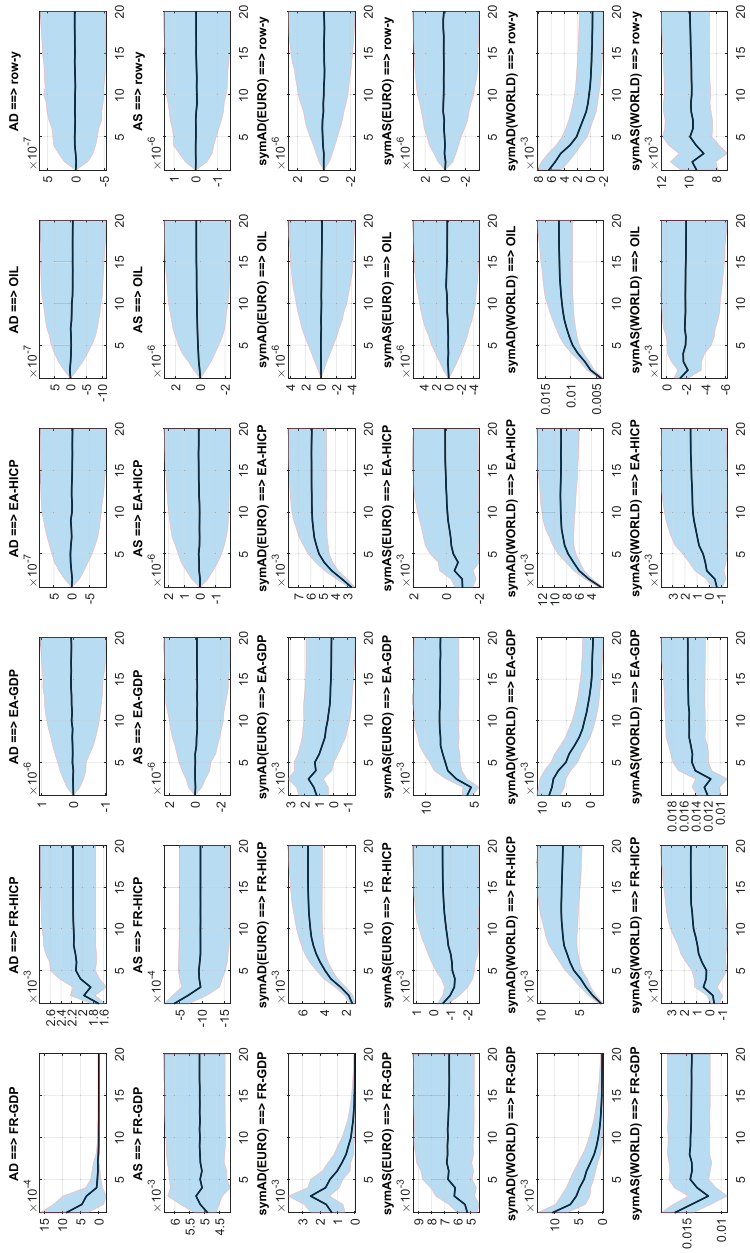


Figure A.3. Impulse Response Functions for Italy (median and 68 percent posterior bands)

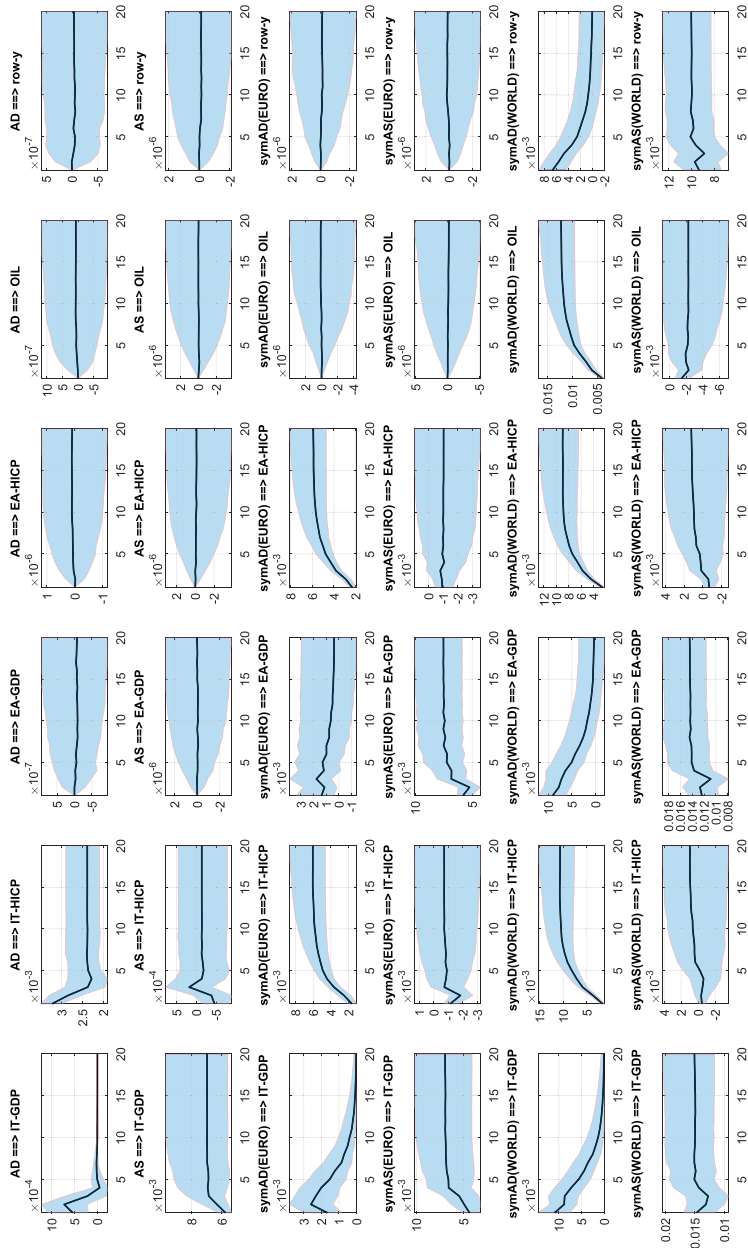


Figure A.4. Impulse Response Functions for Spain
(median and 68 percent posterior bands)

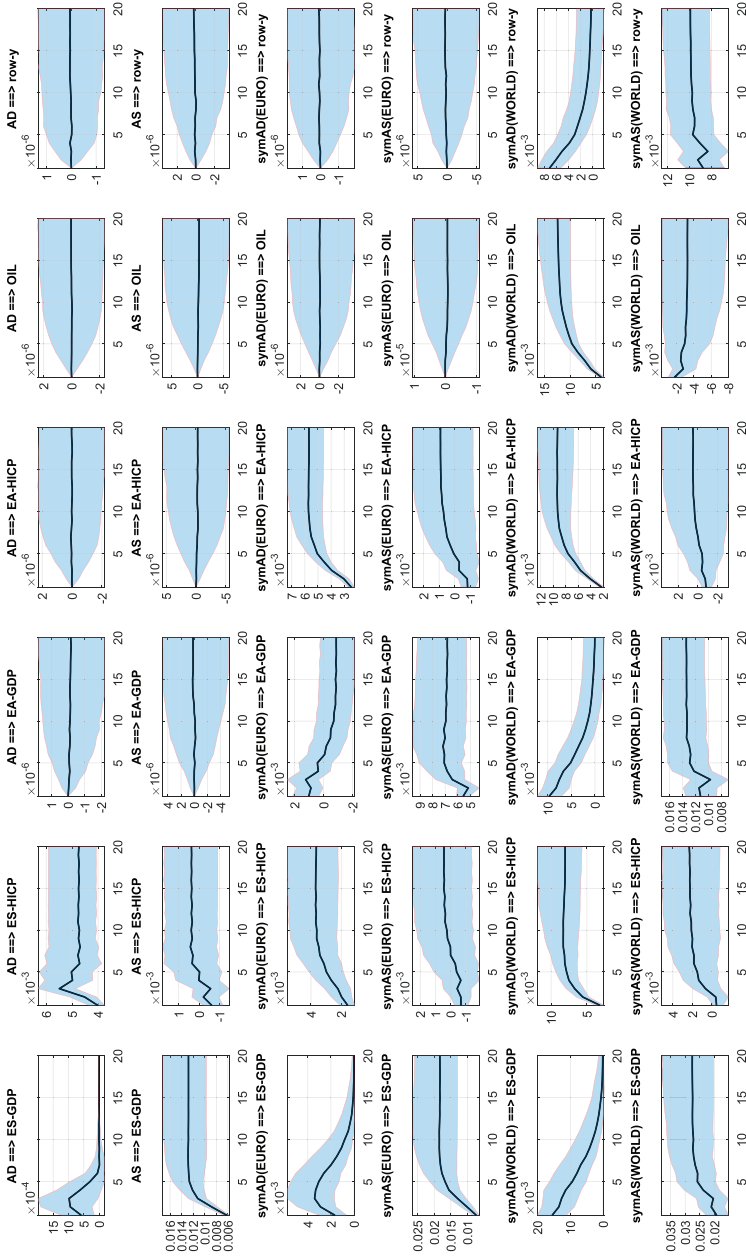
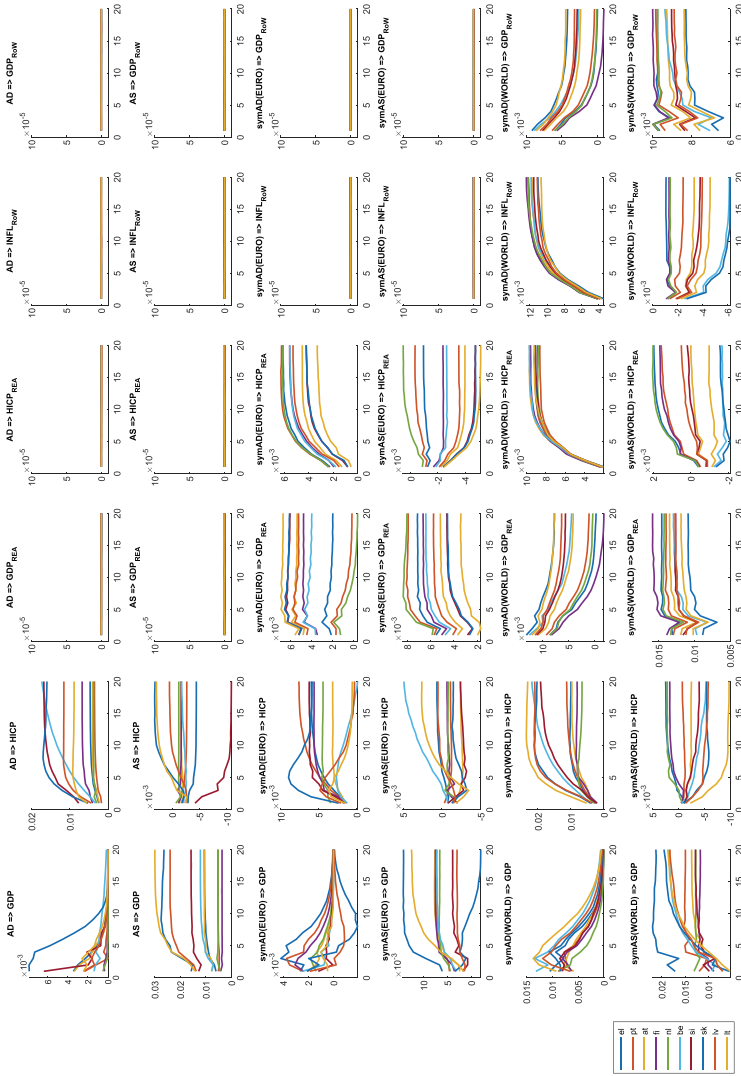


Figure A.5. Impulse Response Functions for Small Countries (medians)



A.3 *Historical Shock Decomposition*

Figure A.6. Germany (median and 68 percent posterior bands)

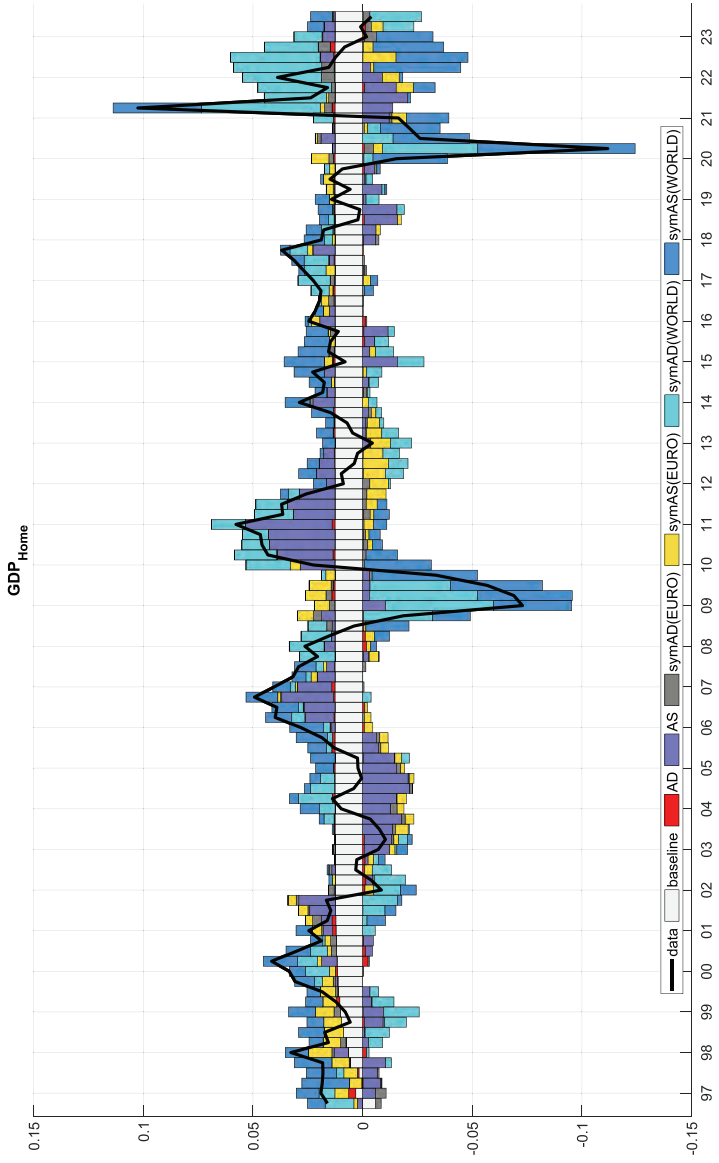


Figure A.7. France (median and 68 percent posterior bands)

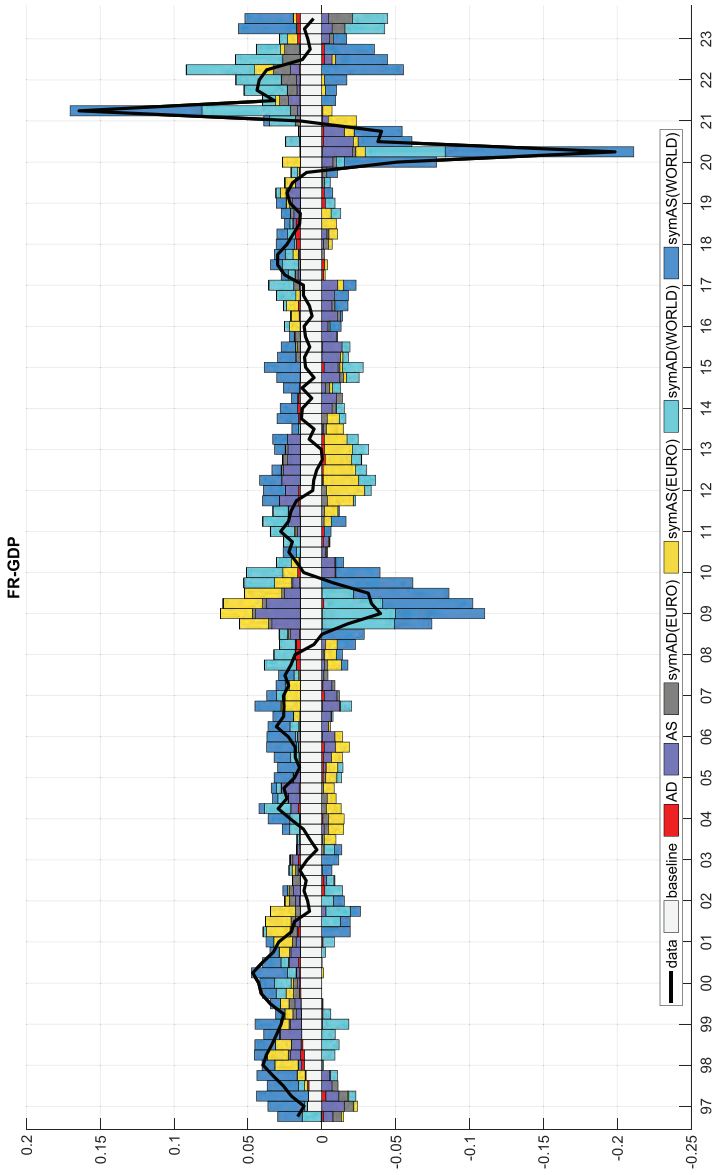


Figure A.8. Italy (median and 68 percent posterior bands)





Figure A.9. Spain (median and 68 percent posterior bands)



A.4 Convergence of Markov Chain Monte Carlo Algorithm

Estimated VARs all converge fast, with posterior statistics all converging already after only 100 or 200 admissible models are found. Our results are all based on overall 1,000 admissible draws. We follow Primiceri (2005) to address convergence of the Gibbs sampling algorithm formally and to assess how efficient the algorithm used explores the posterior. Figures A.10 and A.11 plot the 10th-order sample autocorrelation of the saved draws from the posterior of the reduced-form parameters (regression parameters β and covariance matrix Ω). Figure A.10 shows autocorrelation of the elements of β , and in Figure A.11 are elements of the lower triangular part of the error covariance matrix Ω . Both figures point to weakly autocorrelated MCMC draws—estimated autocorrelation is only rarely outside interval $[-0.05, 0.05]$. Figures A.12 and A.13 plot the inefficiency factors for the posterior estimates of the reduced-form parameters when using a 4 percent tapered window for the estimation of the spectral density at frequency zero. The inefficiency factor¹⁸ is inverse of Geweke (1992) relative numerical efficiency and it serves to quantify the relative efficiency loss in the computation from correlated versus independent samples (Chib 2001). Primiceri (2005) suggests that values of the inefficiency factors below or around 20 would be considered satisfactory. In our application these values are less than two, showing a strong convergence of Gibbs sampler.

¹⁸MCMC draws are realizations of a Markov chain and are, by definition, correlated. When the inefficiency factor is equal to m , we need an MCMC sample that is m times larger compared to an uncorrelated sample to have the same information contained in both.

Figure A.10. Sample Autocorrelation (10th Order) of the Posterior Estimates of the Reduced-Form Parameter Beta

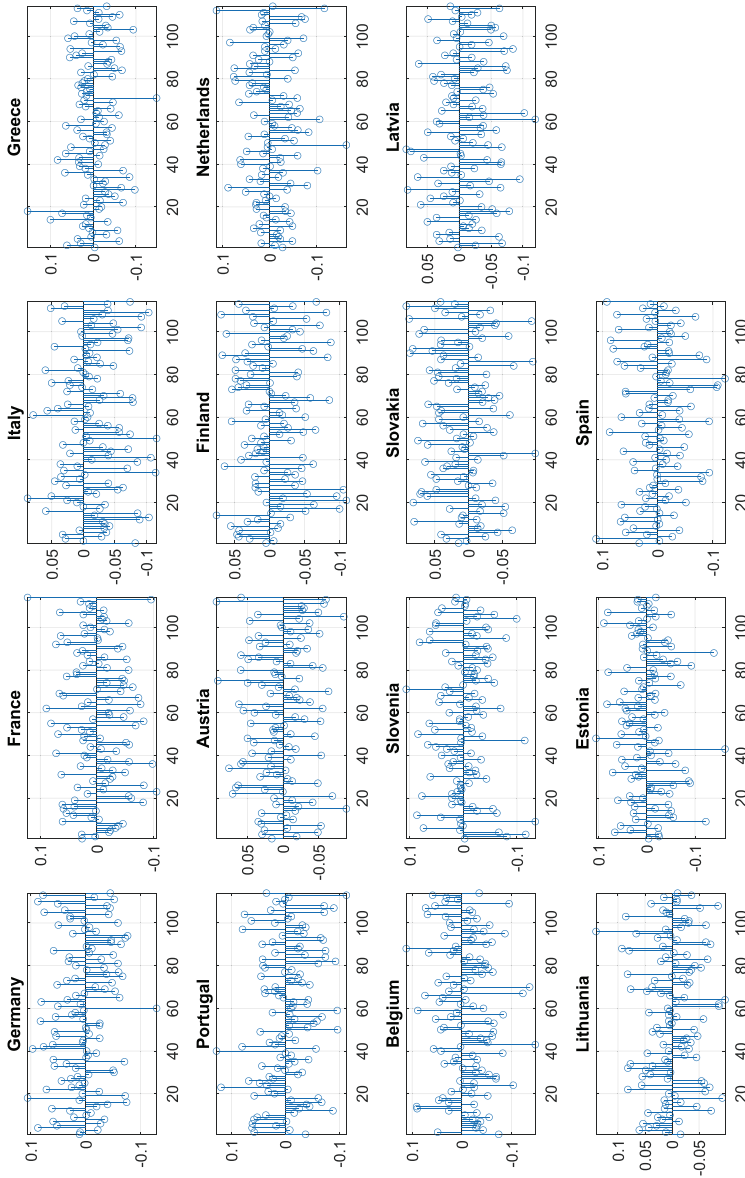


Figure A.11. Sample Autocorrelation (10th Order) of the Posterior Estimates of the Reduced-Form Parameter Omega

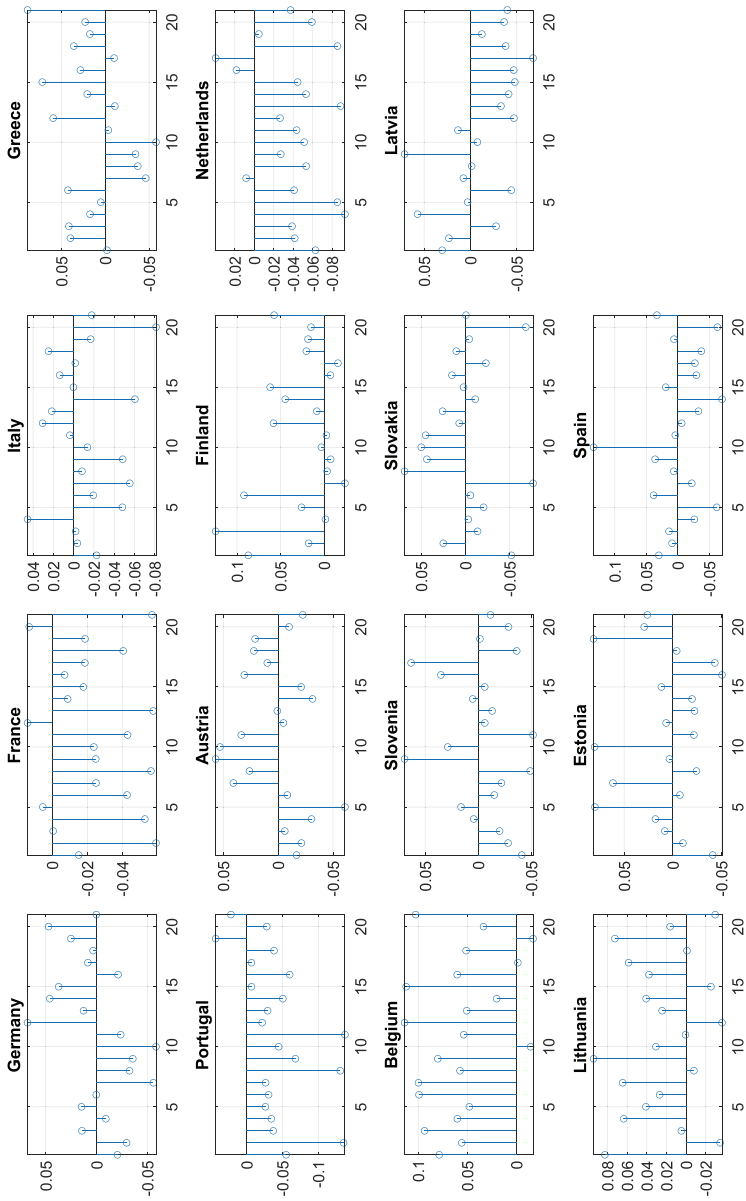


Figure A.12. Inefficiency Factors for the Posterior Estimates of the Reduced-Form Parameter Beta




Figure A.13. Inefficiency Factors for the Posterior Estimates of the Reduced-Form Parameter Omega

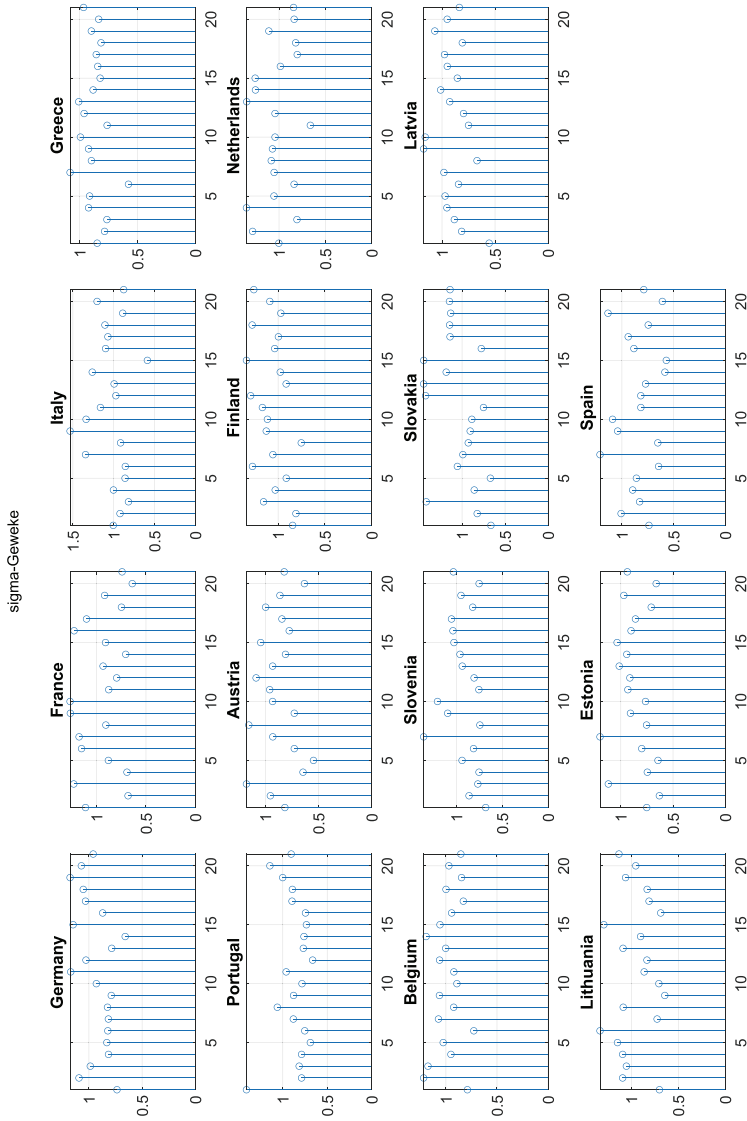


Figure A.14. Output Gap Estimate Based on a BVAR Model: Accumulated Contribution of All Demand Shocks to GDP




Figure A.15. Relative Importance of Local vs. Euro-Area and Global Symmetric Shocks

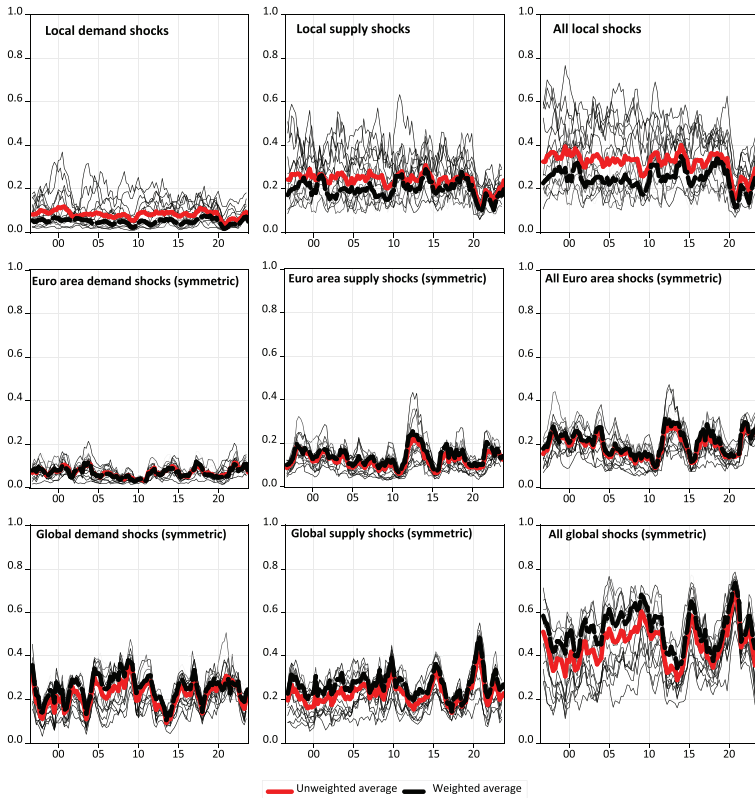


Figure A.16. Decomposition of OCA Index Growth Rates

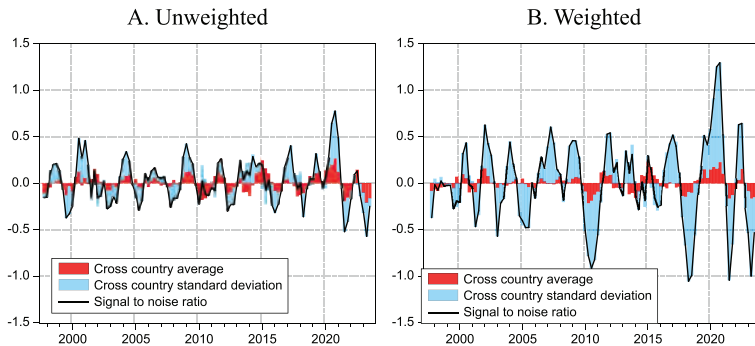


Table A.2. Difference-in-Difference Estimation

Method	All Shocks (AD and AS)		Aggregate Demand (AD)		Aggregate Supply (AS)	
	Coefficient	p-value	Coefficient	p-value	Coefficient	p-value
	<i>All Symmetric Shocks</i>					
Two-Way D-I-D Callaway and Sant'Anna	0.5%	0.86	0.2%	0.91	0.3%	0.84
	4.4%	0.03	2.5%	0.02	1.9%	0.08
	<i>Euro-Area-Specific Symmetric Shocks</i>					
Two-Way D-I-D Callaway and Sant'Anna	-0.8%	0.53	-0.2%	0.84	-0.6%	0.35
	0.7%	0.42	0.5%	0.32	0.2%	0.73
	<i>Global Symmetric Shocks</i>					
Two-Way D-I-D Callaway and Sant'Anna	1.3%	0.46	0.3%	0.68	0.9%	0.38
	3.7%	0.01	2.0%	0.0	1.7%	0.02

Note: AD denotes Aggregate Demand and AS denotes Aggregate Supply. All equations are estimated on 1996:Q3-2022:Q4 for nine countries: Slovenia, Slovakia, Estonia, Latvia, Lithuania, Croatia, Hungary, Poland, and Bulgaria.

References

- Alesina, A., R. J. Barro, and S. Tenreyro. 2002. "Optimal Currency Areas." NBER Working Paper No. 9072.
- Antolin-Diaz, J., and J. Rubio-Ramírez. 2018. "Narrative Sign Restrictions for SVARs." *American Economic Review* 108 (10): 2802–29.
- Arias, J. E., J. F. Rubio-Ramírez, and D. F. Waggoner. 2014. "Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications." International Finance Discussion Paper No. 1100, Board of Governors of the Federal Reserve System.
- . 2018. "Inference Based on Structural Vector Autoregressions Identified with Sign and Zero Restrictions: Theory and Applications." *Econometrica* 86 (2): 685–720.
- Artis, M. J. 2003. "Is There a European Business Cycle?" CESifo Working Paper No. 1053.
- Audzei, V., and F. Brazdik. 2018. "Exchange Rate Dynamics and Their Effect on Macroeconomic Volatility in Selected CEE Countries." *Economic Systems* 42 (4): 584–96.
- Bayoumi, T., and B. Eichengreen. 1992. "Is There a Conflict Between EC Enlargement and European Monetary Unification?" NBER Working Paper No. 3950.
- . 1993. "One Money or Many? On Analyzing the Prospects for Monetary Unification in Various Parts of the World." Working Paper No. 233213, Center for International and Development Economics Research (CIDER), University of California-Berkeley.
- . 1997. "Ever Closer to Heaven: An Optimum-Currency-Area Index for European Countries." *European Economic Review* 41 (3–5): 761–70.
- . 1998. "Exchange Rate Volatility and Intervention: Implications of the Theory of Optimum Currency Areas." *Journal of International Economics* 45 (2): 191–209.
- Beaudry, P., T. J. Carter, and A. Lahiri. 2022. "Looking Through Supply Shocks versus Controlling Inflation Expectations: Understanding the Central Bank Dilemma." Staff Working Paper No. 22-41, Bank of Canada.

- Belke, A., C. Domnick, and D. Gros. 2017. "Business Cycle Synchronization in the EMU: Core vs. Periphery." *Open Economies Review* 28 (5): 863–92.
- Bhatia, R., and C. Davis. 2000. "A Better Bound on the Variance." *American Mathematical Monthly* 107 (4): 353–57.
- Blanchard, O. 2018. "Olivier Blanchard Provides a Brief Reaction to 'Real-Time Estimates of Potential GDP,' by Coibion, Gorodnichenko, and Ulate." Center on Budget and Policy Priorities. <https://www.cbpp.org/research/full-employment/olivier-blanchard-provides-a-brief-reaction-to-real-time-estimates-of>.
- Blanchard, O. J., and D. Quah. 1989. "The Dynamic Effects of Aggregate Demand and Aggregate Supply." *American Economic Review* 79 (4): 655–73.
- Bobeica, E., and M. Jarociński. 2019. "Missing Disinflation and Missing Inflation: A VAR Perspective." *International Journal of Central Banking* 15 (1): 199–232.
- Callaway, B., and P. H. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225 (2): 200–230.
- Camacho, M., G. Perez-Quiros, and L. Saiz. 2006. "Are European Business Cycles Close Enough to be Just One?" *Journal of Economic Dynamics and Control* 30 (9–10): 1687–1706.
- Camba-Mendez, G., and D. Rodriguez-Palenzuela. 2003. "Assessment Criteria for Output Gap Estimates." *Economic Modelling* 20 (3): 529–62.
- Campos, N. F., J. Fidrmuc, and I. Korhonen. 2019. "Business Cycle Synchronisation and Currency Unions: A Review of the Econometric Evidence Using Meta-analysis." *International Review of Financial Analysis* 61 (January): 274–83.
- Campos, N. F., and C. Macchiarelli. 2016. "Core and Periphery in the European Monetary Union: Bayoumi and Eichengreen 25 Years Later." *Economics Letters* 147 (October): 127–30.
- Canova, F., and G. De Nicrolo. 2002. "Monetary Disturbances Matter for Business Fluctuations in the G-7." *Journal of Monetary Economics* 49 (6): 1131–59.
- Caporale, G. M., R. De Santis, and A. Girardi. 2015. "Trade Intensity and Output Synchronisation: On the Endogeneity Properties of EMU." *Journal of Financial Stability* 16 (February): 154–63.

- Carriero, A., T. E. Clark, and M. Marcellino. 2019. "Large Bayesian Vector Autoregressions with Stochastic Volatility and Non-conjugate Priors." *Journal of Econometrics* 212 (1): 137–54.
- Chan, J. C. C. 2019. "Asymmetric Conjugate Priors for Large Bayesian VARs." Working Paper No. 51/2019, Centre for Applied Macroeconomic Analysis (CAMA), Crawford School of Public Policy, The Australian National University.
- Chen, M. J., and L. Gornicka. 2020. "Measuring Output Gap: Is It Worth Your Time?" IMF Working Paper No. 20/24.
- Chib, S. 2001. "Markov Chain Monte Carlo Methods: Computation and Inference." In *Handbook of Econometrics*, Vol. 5, ed. J. J. Heckman and E. Leamer, 3569–3649. North-Holland.
- Coibion, O., Y. Gorodnichenko, and M. Ulate. 2017. "The Cyclical Sensitivity in Estimates of Potential Output." Technical Report, National Bureau of Economic Research.
- Comunale, M., and D. Kunovac. 2017. "Exchange Rate Pass-Through in the Euro Area." ECB Working Paper No. 2003.
- Corden, W. M. 1972. *Monetary Integration*. Essays in International Finance No. 93. Princeton, NJ: Princeton University Press.
- De Grauwe, P. 1996. "Monetary Union and Convergence Economics." *European Economic Review* 40 (3–5): 1091–1101.
- . 2003. "The Euro at Stake? The Monetary Union in an Enlarged Europe." *CESifo Economic Studies* 49 (1): 103–21.
- Deskari-Škrbić, M., K. Kotarac, and D. Kunovac. 2020. "The Third Round of Euro Area Enlargement: Are the Candidates Ready?" *Journal of International Money and Finance* 107 (October): Article 102205.
- . 2020. "Twentieth Anniversary of the Euro: Why Are Some Countries Still Not Willing to Join? Economists View." *Comparative Economic Studies* 62 (2): 242–62.
- European Commission. 1990. "One Market, One Money. An Evaluation of the Potential Benefits and Costs of Forming an Economic and Monetary Union." *European Economy* 44 (October).
- Ferroni, F., and B. Klaus. 2015. "Euro Area Business Cycles in Turbulent Times: Convergence or Decoupling?" *Applied Economics* 47 (34–35): 3791–3815.
- Fidrmuc, J., and I. Korhonen. 2003. "Similarity of Supply and Demand Shocks between the Euro Area and the CEECs." *Economic Systems* 27 (3): 313–34.

- Forbes, K., I. Hjortsoe, and T. Nenova. 2018. "The Shocks Matter: Improving Our Estimates of Exchange Rate Pass-Through." *Journal of International Economics* 114 (September): 255–75.
- Frankel, J. A., and A. K. Rose. 1997. "Is EMU More Justifiable Ex Post Than Ex Ante?" *European Economic Review* 41 (3–5): 753–60.
- . 1998. "The Endogeneity of the Optimum Currency Area Criteria." *Economic Journal* 108 (449): 1009–25.
- Frydrych, J., and S. Burian. 2017. "OCA Indexes and Convergence Process in Europe." *Scientific Annals of Economics and Business* 64 (2): 187–97.
- Furlanetto, F., A. Lepetit, Ø. Robstad, J. Rubio-Ramírez, and P. Ulvedal. 2021. "Estimating Hysteresis Effects." Finance and Economics Discussion Series No. 2021–059.
- Geweke, J. 1992. "Evaluating the Accuracy of Sampling-Based Approaches to the Calculations of Posterior Moments." In *Bayesian Statistics*, Vol. 4, ed. J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, 641–49. Clarendon Press.
- Giannone, D., M. Lenza, and L. Reichlin. 2008. "Business Cycles in the Euro Area." NBER Working Paper No. 14529.
- González-Torres, G., J. E. Gumiel, and B. Szörfi. 2023. "Potential Output in Times of Temporary Supply Shocks." *Economic Bulletin* (European Central Bank) 8/2023: Box 3.
- Horváth, R., and L. Komárek. 2003. "Optimum Currency Area Indices: Evidence from the 1990s." Technical Report, Department of Economics, University of Warwick.
- Inklaar, R., and J. de Haan. 2001. "Is There Really a European Business Cycle? A Comment." *Oxford Economic Papers* 53 (2): 215–20.
- International Monetary Fund. 2011. "Staff Report for the 2011 Article IV Consultation for Germany." Technical Report.
- Ishiyama, Y. 1975. "Theory of Optimum Currency Areas: A Survey." *IMF Staff Papers* 22 (2): 344–83.
- Kadiyala, K. R., and S. Karlsson. 1997. "Numerical Methods for Estimation and Inference in Bayesian VAR-Models." *Journal of Applied Econometrics* 12 (2): 99–132.
- Kenen, P. B. 1969. "The Optimum Currency Area: An Eclectic View." In *Monetary Problems of the International Economy*, ed. R. A. Mundell and A. K. Swoboda. University of Chicago Press.

- Koop, G., and D. Korobilis. 2010. "Bayesian Multivariate Time Series Methods for Empirical Macroeconomics." *Foundations and Trends in Econometrics* 3 (4): 267–358.
- Kose, M. A., C. Otrok, and C. H. Whiteman. 2003. "International Business Cycles: World, Region, and Country-Specific Factors." *American Economic Review* 93 (4): 1216–39.
- Krugman, P. 1993. "Lessons from Massachusetts for EMU." In *Adjustment and Growth in the European Monetary Union*, ed. F. Torres and F. Giavazzi, 241–61. Cambridge University Press.
- Lindeman, R. H., P. F. Merenda, and R. Z. Gold. 1980. *Introduction to Bivariate and Multivariate Analysis*. Glenview, IL: Scott, Foresman and Company.
- Madeira, C., J. Madeira, and P. S. Monteiro. 2023. "The Origins of Monetary Policy Disagreement: The Role of Supply and Demand Shocks." BIS Working Paper No. 1118.
- Masson, P. R., and M. P. Taylor. 1993. *Policy Issues in the Operation of Currency Unions*. Cambridge University Press.
- McKinnon, R. 1963. "Optimum Currency Areas." *American Economic Review* 53: 717–24.
- Mink, M., J. P. Jacobs, and J. de Haan. 2012. "Measuring Coherence of Output Gaps with an Application to the Euro Area." *Oxford Economic Papers* 64 (2): 217–36.
- Miranda-Agrippino, S., and G. Ricco. 2018. "Bayesian Vector Autoregressions." Working Paper No. 756, Bank of England.
- Mongelli, F. P. 2002. "New Views on the Optimum Currency Area Theory: What is EMU Telling Us?" ECB Working Paper No. 138.
- Muggenthaler, P., J. Schroth, and Y. Sun. 2021. "The Heterogeneous Economic Impact of the Pandemic across Euro Area Countries." *Economic Bulletin* (European Central Bank) 5/2021: Box 3.
- Mundell, R. A. 1961. "A Theory of Optimum Currency Areas." *American Economic Review* 51 (4): 657–65.
- . 1975. "Uncommon Arguments for Common Currencies." In *The Economics of Common Currencies*, ed. H. G. Johnson and A. K. Swoboda. London: Allen Unwin Ltd.
- Orphanides, A. 2020. "The ECB's Instruments for Crises and Normal Times: Considerations for the Policy Strategy Review." Working Paper No. 6233-20, MIT Sloan School of Management.

- Peersman, G. 2011. "The Relative Importance of Symmetric and Asymmetric Shocks: The Case of United Kingdom and Euro Area." *Oxford Bulletin of Economics and Statistics* 73 (1): 104–18.
- Primiceri, G. E. 2005. "Time Varying Structural Vector Autoregressions and Monetary Policy." *Review of Economic Studies* 72 (3): 821–52.
- Rubio-Ramírez, J. F., D. F. Waggoner, and T. Zha. 2010. "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference." *Review of Economic Studies* 77 (2): 665–96.
- Sims, C. A., and T. Zha. 1998. "Bayesian Methods for Dynamic Multivariate Models." *International Economic Review* 39 (4): 949–68.
- Skorepa, M. 2013. "Troubles in the Euro Area Periphery: The View through the Lens of a Simple Convergence-Sensitive Optimum Currency Area Index." *Finance a Uver (Czech Journal of Economics and Finance)* 63 (2): 129–51.
- Tenreyro, S., N. Bandera, L. Barnes, M. Chavaz, and L. von dem Berge. 2023. "Monetary Policy in the Face of Supply Shocks: The Role of Inflation Expectations." Paper presented at the ECB Forum on Central Banking, Sintra, Portugal, June 26–28.
- Vieira, C., and I. Vieira. 2012. "Assessing the Endogeneity of OCA Conditions in EMU." *The Manchester School* 80 (1): 77–91.
- Willett, T. D., O. Permpoon, and C. Wihlborg. 2010. "Endogenous OCA Analysis and the Early Euro Experience." *World Economy* 33 (7): 851–72.

Inflation Risk and the Labor Market: Beneath the Surface of a Flat Phillips Curve*

Sirio Aramonte
Federal Reserve Board

While the Phillips curve appeared quiescent after the Great Financial Crisis (GFC), inflation risk, as gauged from option prices, remained sensitive to employment dynamics. Using Phillips-curve regressions centered on option-implied risk-neutral moments, I show that, in tight labor markets, a fall in the unemployment gap raises the risk that inflation overshoots expectations—even if realized inflation remains stable. In tight labor markets, implied moments convey valuable information, as shown by their ability to anticipate future patterns in inflation break-evens and wage growth. Being risk neutral, option-implied moments embed risk premia, which can make moments particularly informative about developments that matter the most to investors. The usefulness of inflation options in assessing risk, despite their illiquidity, is rooted in reputational incentives that dealers have to disseminate accurate quotes.

JEL Codes: G12, G14, G23.

1. Introduction

The weakening empirical link between inflation and the labor market has garnered considerable attention since the Great Financial Crisis (GFC). Two main economic forces stand behind the flattening of the Phillips curve.¹ The first is firmer anchoring of inflation expectations

*I would like to thank Fernando Avalos, Benoît Mojon, Andreas Schrimpf, an anonymous referee, and seminar participants at the Bank for International Settlements. This article represents the views of the author and not those of the Federal Reserve System, its principals, or other members of its staff. Author contact: sirio.aramonte@frb.gov, +1 202-973-6947.

¹The increased responsiveness of monetary policy to job market dynamics can also flatten the slope of the reduced-form Phillips curve, even if the structural

(Ball and Mazumder 2011; Watson 2014; Blanchard 2016; Barnichon and Mesters 2021). The second is the decline in labor bargaining power, which has dampened the prices-wages amplification channel (Lombardi, Riggi, and Viviano 2020; Ratner and Sim 2020; Stansbury and Summers 2020).² Yet, there is evidence that the Phillips curve, while dormant, could resurface rapidly in an overheated job market (Hooper, Mishkin, and Sufi 2020).

In this paper, I study whether perceived upside risk to inflation becomes more sensitive to employment conditions as slack tightens, even if realized inflation does not respond. I do so by characterizing the link between the unemployment gap and option-implied moments of expected U.S. inflation, conditional on labor market tightness. The baseline proxy for tightness is changes in the labor force participation rate, which includes information not necessarily incorporated in the unemployment rate and in the unemployment gap (Erceg and Levin 2014 and Yellen 2014). For instance, the participation rate reflects the flow of discouraged workers in and out of the employment pool. The conclusions are robust to alternative proxies for labor market tightness, including a quadratic specification in the unemployment gap. While the sample excludes the pandemic to avoid the effects of large dislocations in financial markets, the appendix provides selected results from the sample that includes the pandemic.

The main results show that upside risk rises as the unemployment gap falls—but only if the labor market is tight. Importantly, these dynamics hold even as realized inflation remains flat. That is, option-implied inflation moments indicate that, while the Phillips curve may be quiescent, there are important shifts in perceived inflation risk under the surface. These shifts anticipate a number of subsequent developments, as discussed next.

relation is negatively sloped (Eser et al. 2020; Hooper, Mishkin, and Sufi 2020; McLeay and Tenreyro 2020). Additional factors include issues with measuring inflation expectations (Coibion and Gorodnichenko 2015; Coibion, Gorodnichenko, and Kamdar 2018) and the labor slack (Ball and Mazumder 2019).

²These elements are closely related to demographic changes that also affect low-frequency inflation dynamics, such as fluctuations in the working-age population (Juselius and Takáts 2021) and the higher participation rate of older workers (Mojon and Ragot 2019).

I use a two-pronged approach to assess whether the dynamics of option-implied moments convey relevant information. First, I evaluate whether they signal that inflation expectations are becoming unsettled.³ To do so, I consider the comovement between risk-neutral moments and future realized moments of daily break-even rates, which are market-based inflation expectations built from nominal and real government-bond yields. Second, I explore the link between option-implied moments and the future dispersion of wage growth across industries. The rationale is that upside risk to inflation is likely to result in faster wage growth in industries where employees have better bargaining power. Overall, the results indicate that option-implied moments provide useful insights.

The methodology used to extract inflation moments builds on the popular nonparametric approach of Aït-Sahalia and Duarte (2003) and Kitsul and Wright (2013). It is rooted in the work of Breeden and Litzenberger (1978), who derive the return distribution for the asset underlying a set of options from the prices of calls. Importantly, option-implied moments are risk neutral, meaning that they incorporate risk premia. The presence of risk premia, however, is not a hindrance to the analysis. Rather, it means that moments are particularly informative about developments that matter the most to investors and that affect their behavior, chiefly high-inflation states that command larger risk premia (Hilscher, Raviv, and Reis 2026). Elaborating on the arguments set forth by Feldman et al. (2015), Nagel (2016, p. 214) illustrates the usefulness of risk-neutral variables by writing that a “social-welfare maximizing policy should take into account . . . the price that the public is willing to pay to insure against . . . states of the world. Risk-neutral probabilities capture . . . these aspects.”

When studying option-implied inflation distributions, data quality is an important consideration. Available information suggests that these options were relatively liquid between the end of the GFC and the mid-2010s, when trading slowed down considerably and potentially dried up. Subsequently, prices have mostly reflected dealer quotes. Besides theoretical considerations that prices can be informative even in the absence of trading (Milgrom and Stokey

³Note that comparing moments with long-run realized inflation is impractical, since the sample starts just after the GFC.

1982 and Gizatulina and Hellman 2019), conversations with market participants suggest that quotes are useful because they are disseminated, in part, to facilitate the risk management of legacy option holdings by large intermediaries (see Section 2.2 for a detailed discussion). The question of price informativeness is also an empirical one. As discussed above, option-implied moments anticipate future realized moments of break-even rates and cross-industry dispersion in wage increases.

Option-implied moments and probabilities have been used extensively to characterize the behavior of inflation expectations and associated risk premia. The risk of long-lived deflation has been of particular interest to researchers, as in Fleckenstein, Longstaff, and Lustig (2017) and Hilscher, Raviv, and Reis (2026), who develop a method to gauge the risk of very high or very low inflation. Relatedly, Reis (2020) finds that disagreement among market participants, rather than risk premia, is the main determinant of discrepancies between survey-based and market-implied inflation expectations. The drivers of deflation probabilities are explored by Galati, Gorgi, and Zhou (2018), who find evidence of slight unanchoring in the euro area. Eser et al. (2020) highlight that more limited economic slack played an important role in the rightward shift of the euro-area inflation distribution before the COVID-19 pandemic. A separate strand of literature uses information extracted from options to characterize the interplay of inflation and macroeconomic aggregates. For instance, Mertens and Williams (2021) study how the distribution of inflation was affected by the zero lower bound, while Hilscher, Raviv, and Reis (2022) assess the likelihood that inflation can lower real U.S. public debt.

The analysis in this paper is related to research on predicted inflation distributions built using quantile regressions, which are often specified as Phillips curves (Manzan and Zerom 2013). Applications generally focus on the drivers of inflation tails (Busetti, Caivano, and Rodano 2015) and on inflation-at-risk, which quantifies the likelihood that inflation experiences large negative realizations (Banerjee et al. 2020 and López-Salido and Loria 2024). In this literature, Phillips-curve regressions are used to build expected distributions from the historical comovement of inflation with lagged macroeconomic factors. In contrast, inflation options span the full forward-looking distribution on each date,

and I use Phillips-curve regressions to understand the drivers of the risk-neutral distributions of expected inflation. Additionally, the wide set of options' strike prices allows to measure the perceived risk of events—typically rare disasters—that, while not observed in the data, are concerning enough to investors to affect asset prices and trading activity (see, among others, Krasker 1980 and Santa-Clara and Yan 2010 for a general discussion).

In the remainder of the article, Section 2 discusses the methodology to extract risk-neutral inflation distributions. Section 3 computes Phillips curves with realized and risk-neutral expected inflation, while Section 4 studies option-implied moments using Phillips-curve regressions, focusing on the impact of labor market developments on inflation risk. Section 5 evaluates the information content of implied inflation moments and Section 6 concludes.

2. Implied Inflation Distributions

The first step of the analysis consists of extracting risk-neutral distributions from option prices. The data used to build the densities are from Bloomberg and also include inflation swaps to measure point expectations and interest rates swaps as riskless rates. All variables have a five-year horizon and are available each day. As a result, the distributions refer to average annual inflation over the five years following the day in which they are computed.

The methodology used to extract densities is standard in the literature and originates from Breeden and Litzenberger (1978). They link the cumulative probability distribution for the values of the underlying asset (in this case, inflation) to the first derivative of call prices as a function of strike prices.⁴ Option prices are normally interpolated to obtain a dense set of strikes. Doing so typically introduces inflection points that can result in negative probabilities, a sign that the interpolated prices imply arbitrage opportunities. To address this issue, I follow the approach of Ait-Sahalia and Duarte

⁴The set of options includes puts (known as floors) with strikes equal to -1 percent, -0.5 percent, 0 percent, and 0.5 percent, and calls (known as caps) with strikes equal to 1 percent, 1.5 percent, 2 percent, 2.5 percent, 3 percent, 3.5 percent, 4 percent, 4.5 percent, 5 percent, and 6 percent. I convert caps prices into floors prices using the put-call parity (see Mercurio and Zhang 2017).

(2003). They first transform traded prices to satisfy selected slope and convexity restrictions, limiting the incidence of arbitrage. They then use kernel smoothing to obtain a dense set of prices that inherit the favorable properties of the transformed traded prices.

As in Kitsul and Wright (2013), call prices for a given strike can be obtained using the parameters $\{\hat{\beta}_0(k), \hat{\beta}_1(k)\}$ that minimize the following loss function:

$$L = \sum_{i=1}^N [m_i - \beta_0(k) - \beta_1(k)(k_i - k)]^2 \frac{1}{\hat{h}} K\left(\frac{k_i - k}{\hat{h}}\right), \quad (1)$$

where $\{m_i\}_{i=0}^N$ are the transformed prices, N is the number of strikes with a traded price, $K(u) = \exp(-u^2/2) / \sqrt{2\pi}$ is the Gaussian kernel, and \hat{h} is the estimated optimal bandwidth, computed according to Equation (3.23) in Ait-Sahalia and Duarte (2003).

After optimizing Equation (1) for each daily cross-section of traded prices, I compute cumulative probabilities corresponding to a grid of 10,000 equally spaced strikes between the minimum and maximum traded strikes (-1 percent and 6 percent). From this grid, I obtain percentiles that allow me to calculate higher moments robust to outliers (Bowley 1920 and Moors 1988; see Andrade, Ghysels, and Idier 2015 for an application to inflation skewness from survey data). Specifically, option-implied volatility, skewness, and kurtosis are defined as follows:

$$vol_t^{opt} = \frac{\pi_t^{75,opt} - \pi_t^{25,opt}}{\pi_t^{50,opt}}, \quad (2)$$

$$skew_t^{opt} = \frac{\pi_t^{75,opt} + \pi_t^{25,opt} - 2 \cdot \pi_t^{50,opt}}{\pi_t^{75,opt} - \pi_t^{25,opt}}, \quad (3)$$

$$kurt_t^{opt} = \frac{(\pi_t^{87.5,opt} - \pi_t^{62.5,opt}) + (\pi_t^{37.5,opt} - \pi_t^{12.5,opt})}{\pi_t^{75,opt} - \pi_t^{25,opt}}, \quad (4)$$

where $\pi_t^{n,opt}$ is the n^{th} percentile of the risk-neutral distribution of expected inflation on day t . These measures are the variables of interest in the analysis discussed in the remainder of the paper.


2.1 Empirical Properties of Risk-Neutral Inflation Moments

The three moments follow distinct time-series patterns, but they all tend to fluctuate more sharply in the first part of the sample, until U.S. monetary-policy normalization started in 2015 (Figure 1). From then onward, volatility tended to increase in periods of lower risk-neutral inflation expectations, with a particularly pronounced rise in early 2020 at the beginning of the COVID-19 pandemic. Kurtosis tracked risk-neutral inflation expectations quite closely, peaking at the end of the sample in early 2021. In contrast, the behavior of skewness was more nuanced. At first, it turned negative as risk-neutral expected inflation rose in 2018—implying a shift in the probability mass toward moderately higher inflation together with a higher tail risk of markedly lower inflation. However, higher risk-neutral expected inflation in early 2021 was accompanied by a rapid increase in skewness, signaling a higher chance of somewhat lower-than-expected price rises but a more pronounced tail risk of high inflation.

The three moments can be mapped into the probability of various future inflation scenarios. This exercise is especially useful because studying moments allows to understand nuanced changes in the structure of inflation risks, but evaluating the practical relevance of joint dynamics in moments is not immediately intuitive. To gauge upside risk, I consider the link between moments and the likelihood of an inflation “overshoot” (average annual inflation up to 0.5 percentage points above median option-implied inflation) and of an inflation “surge” (more than 2 percentage points above the median). Due to the nonparametric nature of the distributions extracted from Equation (1), I establish the link using regressions robust to outliers (Li 1985) rather than formulas. The dependent variable is the logit transformation of the probability of either scenario, while the independent variables are standardized moments (see Table 1; volatility is replaced by its natural logarithm).

The coefficients shown in the top panel of Table 1 indicate that higher volatility reduces the probability of an overshoot but increases that of a surge. Higher skewness—which reflects a leftward shift in probability mass but a thicker right tail—lowers the chance of an overshoot but heightens the risk of a surge. Kurtosis has a similar effect. The high adjusted R^2 s indicate that the link between

Figure 1. Implied Inflation Moments Over Time



Note: The three panels depict the time series of quarterly averages for option-implied (risk-neutral) inflation volatility, skewness, and kurtosis. Each panel also shows the median of the option-implied distribution in the background for reference. Data are quarterly and cover 2010 to 2021:Q2.

Table 1. Risk-Neutral Moments and Upside Risk to Inflation

	Dependent Variable	
	Prob. Overshoot	Prob. Surge
vol^{opt}	-0.314*** (-20.22)	0.896*** (19.32)
$skew^{opt}$	-0.121*** (-9.00)	0.136*** (3.37)
$kurt^{opt}$	-0.051*** (-2.78)	0.386*** (7.05)
α	-1.286*** (-99.67)	-3.339*** (-86.56)
Obs.	119	119
Adj. R ²	0.829	0.791
Probability Values		
	Prob. Overshoot	Prob. Surge
Prob. if All Moments at Mean	21.7%	3.4%
Δ Prob. if vol^{opt} is 1 σ Above Mean	-4.9%	4.6%
Δ Prob. if $skew^{opt}$ is 1 σ Above Mean	-2.0%	0.5%
Δ Prob. if $kurt^{opt}$ is 1 σ Above Mean	-0.9%	1.5%

Note: The table links risk-neutral moments to the likelihood that future inflation exceeds expectations by a moderate or substantial amount. Moments and probabilities are contemporaneous. In the upper panel, the variable of interest is the logit transformation ($y = \ln\left(\frac{p}{1-p}\right)$) of the probability that annual average realized inflation in the next five years is up to 0.5 percentage points above median expected inflation (“Prob. Overshoot”) or 2 percentage points or more above the median (“Prob. Surge”). The dependent variables are linked to standardized inflation moments using regressions robust to outliers (Li 1985). In the lower panel, the probabilities in the first line are computed with the inverse logit transformation applied to the coefficients α in the upper panel: $(1 + \exp(-\alpha))^{-1}$. Probability changes (Δ Prob) when moments are one standard deviation above the mean are computed as Δ Prob = $(1 + \exp(-(\alpha + mom^{opt})))^{-1} - (1 + \exp(-\alpha))^{-1}$, where mom^{opt} is one of the three option-implied moments. The sample covers 2010–19.

moments and probabilities is tight. When the moments are at their sample average, the probability of an overshoot is 21.7 percent and that of a surge is 3.4 percent (lower panel). One-standard-deviation increases in individual moments alter the two probabilities meaningfully. Changes in skewness affect the chance of an overshoot the

most (−2.0 percentage points), while kurtosis has a larger effect on the likelihood of an inflation surge (+1.5 percentage points).

2.2 How Informative Are the Prices of Inflation Options?

As discussed in Kitsul and Wright (2013), the market for U.S. inflation options developed after the GFC. In 2011, trading amounted to \$22 billion in notional value, which represented a threefold increase relative to 2010 but was still much less than the total for inflation-indexed bonds. Kitsul and Wright (2013) argue that large volumes are not crucial to price formation, pointing to the work Wolfers and Zitzewitz (2004) on markets characterized by small transactions. From a theoretical standpoint, the no-trade theorem illustrates how quoted prices can be informative even when trading does not occur at all (Milgrom and Stokey 1982 and Gizatulina and Hellman 2019).

While details on recent trading activity for inflation options are difficult to obtain, conversations with data providers and market participants paint the picture of instruments with very limited volume but with meaningful legacy open positions. Although the number of transactions fell rapidly around 2015, large financial institutions, reportedly including banks and insurance companies, hold inflation options that will not expire until at least the mid-2020s. The reason is that these instruments have always had long maturities, mostly expiring after 5 or 10 years.

Depending on the accounting and regulatory treatment of such positions, inflation options need to be marked-to-market regularly, possibly every day. Marking-to-market of illiquid instruments is a common industry practice that often relies on dealer quotes. A key reason why dealers provide frequent quotes for instruments that rarely trade is precisely to ensure that clients can value their positions daily. Quotes are often based on the prices of more liquid contracts that load on similar risks—such as interest rate derivatives in the case of inflation options—and on expert judgment.

Dealers that provide quotes for illiquid instruments have two main incentives to give unbiased estimates. The first is that these quotes are actionable, meaning that clients can ask to trade on them. While there is no obligation for a dealer to transact, a large discrepancy between the quote and the proposed trade price would be detrimental to the dealer's reputation. The second incentive relates

to the use of these quotes. To the extent that they are inputs to the calculation of regulatory capital—as they reportedly can be—they would be validated in a number of ways to ensure that they are of appropriately good quality. This process can include, among other steps, a comparison with available transaction prices and with consensus pricing data that aggregates estimates from a large set of market participants. Providing quotes that do not clear these hurdles would tarnish the reputation of a dealer and affect its business.

The appendix includes an empirical analysis of the informativeness of option-implied distributions.

3. Phillips Curves with Realized and Implied Inflation

Before exploring the link between risk-neutral inflation moments and the labor market, I estimate standard Phillips curves with realized inflation to provide an initial reference point. The specification is similar to the one in Hooper, Mishkin, and Sufi (2020), who measure inflation with the changes in the core personal consumption expenditures (PCE) index. The only difference is that I use changes in the consumer price index (CPI), since the payoff of U.S. inflation options is tied to CPI. The baseline specification, using data at the quarterly frequency, is

$$\pi_t = \alpha + \beta \cdot ugap_t + \sum_{j=1}^3 \eta_j \pi_{t-j} + \theta \cdot \pi_{t-1}^{e,frb} + \lambda \cdot \Delta RelImport_{t-1} + \epsilon_t, \quad (5)$$

where π is the period-on-period relative change in CPI, $ugap$ is the difference between the unemployment rate and the natural rate of unemployment, $\pi^{e,frb}$ is the inflation expectation series used in the Federal Reserve's FRB/US model, and $\Delta RelImport$ is the change in the relative price of import goods relative to domestic goods. Standard errors are based on Newey and West (1987) with four lags. The variable $\pi^{e,frb}$ is available from the Federal Reserve Board, while all others are from the Federal Reserve Economic Data (FRED) database of the Federal Reserve Bank of St. Louis. Regression coefficients are computed using data from 2010 to 2019, excluding the last 18 months of the sample to avoid the influence of the large dislocations

caused by the COVID-19 pandemic (see the appendix for selected results based on an extended sample).

The results in the first column of Table 2 can be compared with those in Table 2.5 of Hooper, Mishkin, and Sufi (2020), since the sample is nearly identical. The conclusions are quite similar, in that the slope of the Phillips curve relative to the unemployment gap is negative but not statistically significant, implying that the measured Phillips curve is flat. Early studies, such as Tobin (1972), recognized that this relation could be asymmetric and depend on the state of the labor market. As a result, the second column of the table shows coefficients from a specification where the effect of the unemployment gap is conditioned on the labor market. The conditioning variable is the 12-month trailing change in the labor force participation rate ($\Delta part$), which, as discussed in Section 1, is a broad measure of labor slack that reflects the flow of discouraged workers in and out of the labor pool. The regression is

$$\pi_t = \alpha + \beta \cdot ugap_t + \gamma \cdot \Delta part_{t-1} + \delta \cdot ugap_t \cdot \Delta part_{t-1} + \sum_{j=1}^3 \eta_j \pi_{t-j} + \theta \cdot \pi_{t-1}^{e,frb} + \lambda \cdot \Delta RelImport_{t-1} + \epsilon_t. \quad (6)$$

The coefficients are quite similar to those in the unconditional specification. In addition, the interaction term is not statistically significant, meaning that the marginal effect of the unemployment gap on inflation is statistically zero in both slack and tight labor markets (bottom panel of the table). In the two middle columns of Table 2, the sample focuses on the post-GFC period, for which inflation-option data are available. Once again, the Phillips curve is flat. In addition, the interaction term with the unemployment gap is not statistically significant, meaning that the slope of the Phillips curve does not depend on labor slack.

The positive coefficient on $\Delta part$ is consistent with the participation rate reflecting inflationary increases in labor demand rather than deflationary improvements in labor supply. Recent studies indicate that labor demand can raise the participation rate, especially for under-represented workers (Hobijn and Şahin 2021), even if labor supply shapes participation in subsets of the population (Mojon and Ragot 2019). Yellen (2014, p. 5) highlighted the role of labor demand

Table 2. Phillips Curve with Realized and Implied Inflation (Quarterly)

Period (Quarterly) →	1989–2019		2010–19		2010–19	
	π_t		π_t		$\pi_t^{e,opt}$	
Dep. Var. →						
$ugapt$	-0.180 (-1.54)	-0.279 (-1.22)	-0.144 (-0.61)	0.409 (1.50)	0.281* (1.91)	0.590*** (2.80)
$\Delta part_{t-1}$		-0.070 (-0.54)		0.585*** (4.00)		0.409*** (3.12)
$ugapt \cdot \Delta part_{t-1}$		-0.066 (-0.70)		-0.085 (-0.74)		0.108 (1.37)
π_{t-1}	-0.120 (-0.88)	-0.128 (-0.92)	0.113 (0.55)	0.086 (0.47)	0.473*** (2.93)	0.403** (2.48)
π_{t-2}	-0.220 (-0.91)	-0.211 (-0.92)	-0.041 (-0.20)	-0.034 (-0.17)	0.618** (2.22)	0.497* (1.90)
π_{t-3}	0.019 (0.22)	0.007 (0.08)	-0.340 (-1.66)	-0.271 (-1.55)	-0.028 (-0.11)	-0.029 (-0.11)
$\pi_{t-1}^{e,frb}$	0.411*** (2.68)	0.446** (2.36)	0.312 (1.64)	0.195 (1.25)	-0.277 (-1.24)	-0.314 (-1.47)
$\pi_{t-1}^{e,opt}$					0.618*** (5.18)	0.759*** (6.14)
$\Delta ReImport_{t-1}$	0.167 (0.92)	0.155 (0.88)	0.176 (0.88)	0.112 (0.50)	-0.310 (-1.30)	-0.350 (-1.33)
Obs.	123	123	40	40	39	39
R ²	0.117	0.124	0.132	0.266	0.621	0.667

(continued)

Table 2. (Continued)

Labor Market	Marginal Effect of $ugap_t$	
Slack ($\Delta part_{t-1} = -1\sigma$)	-0.213 (-1.26)	0.494 (1.57)
Tight ($\Delta part_{t-1} = +1\sigma$)	-0.345 (-1.13)	0.699*** (3.14)

Note: The table shows coefficients from Phillips curves similar to those in Hooper, Mishkin, and Sufi (2020), using quarterly data. When the dependent variable is realized inflation (π_t), the specification is $\pi_t = \alpha + \beta \cdot ugap_t + \gamma \cdot \Delta part_{t-1} + \delta \cdot ugap_t \cdot \Delta part_{t-1} + \sum_{j=1}^3 \eta_j \pi_{t-j} + \theta \cdot \pi_{t-1}^{e,frb} + \lambda \cdot \Delta RelImport_{t-1} + \epsilon_t$. When the dependent variable is option-implied expected inflation ($\pi_t^{e,opt}$), the specification is $\pi_t^{e,opt} = \alpha + \beta \cdot ugap_t + \gamma \cdot \Delta part_{t-1} + \delta \cdot ugap_t \cdot \Delta part_{t-1} + \sum_{j=1}^3 \eta_j \pi_{t-j} + \theta \cdot \pi_{t-1}^{e,frb} + \lambda \cdot \Delta RelImport_{t-1} + \epsilon_t$. The variables included in the regressions are the relative period-on-period change in the consumer price index (π), median expected inflation as implied from option prices ($\pi^{e,opt}$), the difference between the unemployment rate and the natural rate of unemployment ($ugap$), the four-quarter change in the labor force participation rate ($\Delta part$), the inflation expectation series used in the Federal Reserve's FRB/US model ($\pi^{e,frb}$), and the change in the relative price of import goods relative to domestic goods ($\Delta RelImport$). All variables are standardized and constants are unreported. t -statistics are calculated with standard errors based on Newey and West (1987) with four lags.

as a driving factor, writing that then-recent dynamics in “labor force participation . . . could partly reflect discouraged workers rejoining the labor force in response to the significant improvements . . . in labor market conditions.”

The closest specification to a Phillips curve that uses information from inflation derivatives is a regression of risk-neutral expected inflation on the unemployment gap and standard controls. Since these expectations are risk neutral, they embed inflation risk premia, which could themselves be linked to the state of the labor market. The last two columns of Table 2 show coefficients from Equations (5) and (6), respectively, where the left-hand variable is $\pi^{e,opt}$ instead of realized inflation. This variable is the median of the risk-neutral distribution of expected inflation extracted from option prices. Since inflation swaps are used to convert the prices of floors into those of caps, $\pi^{e,opt}$ closely tracks the inflation swap rate. The regressions include lagged $\pi^{e,opt}$ as a market-based counterpart to $\pi^{e,frb}$. The results indicate that the Phillips curve is *positively* sloped (upper panel), a finding that I discuss in detail in the next section. The slope is considerably larger in tight labor markets (lower panel).

As shown in Table 3, this pattern also holds with monthly rather than quarterly data. In addition, coefficients on lagged realized and risk-neutral expected inflation indicate short-term persistence in both variables. The monthly specification, reported below, does not include $\Delta RelImport_{t-1}$ as a regressor due to the higher data frequency:

$$\begin{aligned} \pi_t = & \alpha + \beta \cdot ugap_t + \gamma \cdot \Delta part_{t-1} + \delta \cdot ugap_t \cdot \Delta part_{t-1} \\ & + \sum_{j=1}^3 \eta_j \pi_{t-j} + \zeta \cdot \pi_{t-1}^{e,opt} + \epsilon_t. \end{aligned} \quad (7)$$

3.1 The Role of Monetary Policy Expectations

The observation that option-implied expected inflation rises when the unemployment gap widens appears counterintuitive, since Phillips curves normally have a negative slope. However, market-based inflation expectations, on which I focus in this part of the analysis, are different in key respects from realized inflation contemporaneous to the unemployment gap. Besides incorporating risk

Table 3. Phillips Curve with Realized and Implied Inflation (Monthly)

Period (Monthly) →	2010–19		2010–19	
Dep. Var. →	π_t		$\pi_t^{e,opt}$	
$ugap_t$	0.025 (0.41)	0.101 (1.10)	0.036 (0.98)	0.156** (2.47)
$\Delta part_{t-1}$		0.105 (1.18)		0.150** (2.38)
$ugap_t \cdot \Delta part_{t-1}$		-0.039 (-0.64)		0.059 (1.43)
π_{t-1}	0.429*** (4.88)	0.421*** (4.80)	0.182*** (3.31)	0.184*** (3.23)
π_{t-2}	-0.131* (-1.74)	-0.129* (-1.75)	-0.048 (-1.11)	-0.038 (-0.83)
π_{t-3}	-0.060 (-0.66)	-0.065 (-0.73)	-0.033 (-0.53)	-0.034 (-0.54)
$\pi_{t-1}^{e,opt}$	-0.008 (-0.16)	0.007 (0.15)	0.885*** (20.16)	0.902*** (24.20)
Obs.	119	119	119	119
R ²	0.168	0.180	0.836	0.845
Labor Market	Marginal Effect of $ugap_t$			
Slack ($\Delta part_{t-1} = -1\sigma$)	0.140 (1.10)		0.097 (1.41)	
Tight ($\Delta part_{t-1} = +1\sigma$)	0.061 (0.67)		0.215*** (2.64)	
<p>Note: The table shows coefficients on Phillips curves similar to those in Hooper, Mishkin, and Sufi (2020), using monthly data. The dependent variable y is either π or $\pi^{e,opt}$. The specification is $y_t = \alpha + \beta \cdot ugap_t + \gamma \cdot \Delta part_{t-1} + \delta \cdot ugap_t \cdot \Delta part_{t-1} + \sum_{j=1}^3 \eta_j \pi_{t-j} + \zeta \cdot \pi_{t-1}^{e,opt} + \epsilon_t$. The variables included in the specifications are defined in Table 2. Monthly natural rates are linearly interpolated from quarterly data. Constants are not reported. All variables are standardized. t-statistics are calculated with standard errors based on Newey and West (1987) with 12 lags.</p>				

premia, expectations refer to inflation that investors anticipate (in a risk-neutral sense) will be realized over the following five years. As a result, in linking current labor market conditions to expectations of future inflation, one needs to consider the effect of possible policy actions in response to the state of the labor market. Naturally, these

actions affect future inflation, as well as its current expectations, much more than they affect contemporaneous realized inflation. Can changes in expected monetary policy contribute to the positive coefficients on the unemployment gap in Tables 2 and 3?

In order to avoid endogeneity issues, I focus on an event study that, in the spirit of Phillips-curve regressions focused on inflation expectations, links inflation expectations to the state of the labor market while controlling for monetary policy expectations. Given the event-study setup, the variables are expressed in changes, except for the labor market variable, which, to address endogeneity, is expressed as surprises relative to consensus surveys. Specifically, the event study focuses on days when the unemployment rate is announced as part of scheduled Employment Situation report releases. Using the specification below, I condition the link between market-based inflation expectations and unemployment surprises on both expected monetary policy and the labor market:

$$\begin{aligned} \Delta\pi_t^{e,opt} = & \Delta vix_t + surp_t + \Delta part_{t-1} + \Delta ffr_t \\ & + surp_t \cdot \Delta part_{t-1} + surp_t \cdot \Delta ffr_t + \Delta part_{t-1} \cdot \Delta ffr_t \\ & + surp_t \cdot \Delta part_{t-1} \cdot \Delta ffr_t + \epsilon_t, \end{aligned} \quad (8)$$

where Δffr is the change in the expected federal funds rate one year ahead, and $surp$ is the difference between the announced unemployment rate and survey expectations from Thomson Reuters, divided by the natural rate of unemployment (linearly interpolated as needed). As in the rest of the paper, regressions are based on standardized variables, so that the coefficients can be interpreted as (i) the effect of a one-standard-deviation change in the independent variables (ii) in terms of standard deviations of the dependent variable.

As shown in panel A of Table 4, Δffr is an important explanatory variable for $\Delta\pi^{e,opt}$. Its interaction with $\Delta part$ is also negative and strongly statistically significant. The coefficients on $surp$ and on its interaction with $\Delta part$ are not statistically significant. However, given the size and the strong statistical significance of the other interactions, the marginal effect of a wider unemployment gap on option-implied expected inflation is positive and larger in tight labor markets (as in Tables 2 and 3, even if the test designs are completely different)—but only when investors expect looser monetary policy

Table 4. Expected Policy Rates and Implied-Inflation/Unemployment-Gap Link

A. Changes in Implied Inflation and <i>ugap</i> Surprises				
Dep. Var. →	$\Delta\pi_t^{e,opt}$	$\Delta\pi_t^{e,opt}$	$\Delta\pi_t^{e,opt}$	$\Delta\pi_t^{e,opt}$
Δvix_t				-0.142 (-1.37)
$surp_t$	0.158 (1.33)	0.129 (1.33)	0.075 (0.79)	0.083 (0.96)
$\Delta part_{t-1}$	0.170** (2.57)	0.122* (1.93)	0.122 (1.57)	0.111 (1.44)
Δffr_t		0.267* (1.87)	0.348*** (3.95)	0.323*** (3.66)
$surp_t \cdot \Delta part_{t-1}$			0.066 (1.17)	0.048 (0.80)
$surp_t \cdot \Delta ffr_t$			-0.086* (-1.67)	-0.097* (-1.79)
$\Delta part_{t-1} \cdot \Delta ffr_t$			-0.248*** (-2.91)	-0.252*** (-2.93)
$surp_t \cdot \Delta part_{t-1} \cdot \Delta ffr_t$			-0.009 (-0.12)	-0.030 (-0.39)
Obs.	116	116	116	116
R ²	0.057	0.125	0.215	0.234
B. Effect of <i>ugap</i> Surprise on Change in Implied Inflation, by Labor Tightness and Change in Expected Monetary Policy				
	Expected Monetary Policy			
	Expand		Contract	
Labor Market	$(\Delta ffr_t = -1\sigma)$		$(\Delta ffr_t = +1\sigma)$	
Slack $(\Delta part_{t-1} = -1\sigma)$	0.103 (1.65)		-0.032 (-0.14)	
Tight $(\Delta part_{t-1} = +1\sigma)$	0.258** (2.51)		0.004 (0.04)	
<p>Note: The table shows the result of an event study based on scheduled releases of the Employment Situation report. The variables Δvix_t and Δffr_t are changes in log-VIX and in the one-year-ahead federal funds rate (implied from futures) between day $t - 1$ and day t, where t is the date of the monthly release of the Employment Situation report. The variable $surp_t$ is the difference between the actual and expected unemployment rate (based on Thomson Reuters surveys) divided by the natural rate of unemployment measured in the previous month. Monthly natural rates are linearly interpolated from quarterly data. Other variables are defined in Table 2. Constants are not reported. All variables are standardized. t-statistics are calculated with standard errors based on Newey and West (1987) with 12 lags.</p>				

(panel B).⁵ This result confirms that policy actions are key determinants of the empirical link between inflation and the unemployment rate (McLeay and Tenreyro 2020).

4. Inflation Risk and the Labor Market

The Phillips-curve framework can be used to explore how the whole risk-neutral distribution of expected inflation relates to labor dynamics. Using option-implied moments as dependent variables, Phillips-curve regressions can link the state of the job market to different types of inflation risks (and embedded risk premia), from the dispersion of future changes to the incidence of extreme realizations. In turn, such risks can be expressed as the probabilities of relevant inflation scenarios (see Section 2.1). Ultimately, this analysis can shed light on the ebbs and flows of perceived inflationary pressures as employment opportunities change—even when realized inflation remains stable.

In practice, it is reasonable to expect that, as the labor market tightens, inflation risk changes from the possibility of relatively minor movements around expected inflation to a higher likelihood of unusually high inflation. In terms of inflation moments, this change corresponds to lower volatility (inflation risk becomes less symmetric), lower skewness (the probability mass moves toward higher inflation values), and higher kurtosis (more pronounced tail risk).

The empirical analysis builds on specifications similar to the Phillips curves used to study realized inflation at the quarterly and monthly frequencies, as detailed in Equations (6) and (7). Apart from the dependent variable (y_t) being one of volatility, skewness, or kurtosis, the only difference is that the set of regressors includes lagged values of y_t to account for autocorrelation in moments. At the quarterly frequency, the specification features the inflation-expectations series used in the Federal Reserve's FRB/US model

⁵The surprise variable has a standard deviation of 3 percent over the regression sample. Over the same period, daily changes in risk-neutral inflation expectations have a standard deviation of 3.91 basis points. Given the 0.258 coefficient corresponding to expansionary monetary policy expectations and tight labor markets, a one-standard-deviation surprise raises implied inflation by 1 basis point.

($\pi^{e,frb}$) and the change in the relative price of import goods relative to domestic goods ($\Delta RelImport$):

$$y_t = \alpha + \beta \cdot ugap_t + \gamma \cdot \Delta part_{t-1} + \delta \cdot ugap_t \cdot \Delta part_{t-1} + \zeta y_{t-1} + \sum_{j=1}^3 \eta_j \pi_{t-j} + \theta \cdot \pi_{t-1}^{e,frb} + \lambda \cdot \Delta RelImport_{t-1} + \epsilon_t. \quad (9)$$

At the monthly frequency, constraints on data availability mean that inflation expectations are measured with the median of the option-implied distribution ($\pi^{e,opt}$) and that $\Delta RelImport$ is not included:

$$y_t = \alpha + \beta \cdot ugap_t + \gamma \cdot \Delta part_{t-1} + \delta \cdot ugap_t \cdot \Delta part_{t-1} + \zeta y_{t-1} + \sum_{j=1}^3 \eta_j \pi_{t-j} + \zeta \cdot \pi_{t-1}^{e,opt} + \epsilon_t. \quad (10)$$

The results shown in Table 5 clearly indicate that the higher moments respond to the unemployment gap. For skewness, the coefficient on $ugap$ is positive, as is the one on the interaction between $ugap$ and $\Delta part$. The bottom panel shows marginal coefficients for $ugap$ when the labor market is slack or tight, obtained by combining the coefficients on $ugap$ and on the $ugap$ - $\Delta part$ interaction. These marginal coefficients indicate that a smaller unemployment gap is accompanied by lower skewness, that is by a shift in probability mass towards higher inflation—but only if the labor market is tight. For kurtosis, both the main and interaction coefficients are negative, as are the marginal coefficients corresponding to a tight labor market (bottom panel of the table). As a result, when the job market is already tight, a lower unemployment gap increases kurtosis and tail thickness. Volatility also declines as the unemployment gap shrinks, but the marginal coefficients are only statistically significant when using monthly data.

The three moments can be translated into probabilities of future inflation scenarios. Doing so is useful because individual moments can measure nuanced changes in the structure of inflation risks, but assessing the practical relevance of joint moment dynamics is not intuitive.

To better gauge the practical implications of these changes in implied moments, it is useful to translate these changes to changes

Table 5. Phillips-Curve Regressions with Option-Implied Inflation Moments

Period →	2010–19 (Quarterly)			2010–19 (Monthly)		
	vol_t^{opt}	$skew_t^{opt}$	$kurt_t^{opt}$	vol_t^{opt}	$skew_t^{opt}$	$kurt_t^{opt}$
$ugapt$	0.177 (0.79)	0.696*** (2.90)	-0.391* (-1.84)	0.184*** (2.77)	0.342*** (2.72)	-0.231 (-1.63)
$\Delta part_{t-1}$	0.036 (0.28)	0.463*** (3.23)	-0.242 (-1.57)	-0.021 (-0.55)	0.277** (2.36)	-0.198** (-2.57)
$ugapt \cdot \Delta part_{t-1}$	0.103 (1.36)	0.684*** (4.54)	-0.439*** (-3.34)	0.024 (0.64)	0.330*** (3.06)	-0.262*** (-3.68)
y_{t-1}	0.652*** (4.82)	0.030 (0.18)	0.145 (0.54)	0.774*** (11.84)	0.298*** (2.86)	0.377*** (3.90)
π_{t-1}	-0.091 (-0.73)	-0.376** (-2.09)	-0.005 (-0.03)	-0.036 (-0.85)	0.050 (0.44)	0.030 (0.46)
π_{t-2}	-0.412*** (-2.62)	0.516 (1.28)	-0.093 (-0.31)	0.092** (2.18)	-0.031 (-0.22)	-0.168* (-1.88)
π_{t-3}	-0.138 (-1.19)	-0.141 (-0.74)	0.018 (0.10)	0.039** (2.00)	-0.037 (-0.53)	-0.169** (-2.48)
$\pi_{t-1}^{e,frb}$	0.256*** (2.25)	-0.066 (-0.29)	-2.467 (-1.66)	-0.140*** (-2.99)	0.090 (1.12)	0.251*** (3.33)
$\Delta RelImport_{t-1}$	0.066 (0.39)	-0.335 (-0.99)	0.272 (1.00)			
Obs. R ²	39 0.834	39 0.511	39 0.471	$\pi_{t-1}^{e,opt}$ Obs. R ²	119 0.253	119 0.385
Labor Market	Marginal Effect of $ugapt$					
Slack ($\Delta part_{t-1} = -1\sigma$)	0.073 (0.37)	0.012 (0.04)	0.048 (0.33)		0.160** (2.14)	0.031 (0.19)
Tight ($\Delta part_{t-1} = +1\sigma$)	0.280 (1.03)	1.380*** (4.83)	-0.831** (-2.58)		0.208*** (2.71)	-0.493*** (-3.16)

Note: The table shows coefficients from regressions similar to Phillips curves, where the dependent variables are option-implied moments (log-volatility, skewness, and kurtosis) for expected average inflation five years ahead. The variables included in the specifications are defined in Tables 2 and 3. Monthly natural rates are linearly interpolated from quarterly data. Constants are not reported. All variables are standardized. t -statistics are calculated with standard errors based on Newey and West (1987) with 12 lags. The sample covers 2010–19.

in the probabilities of different inflation scenarios. As discussed in Section 2.1, there is a tight empirical relation between implied moments and the probability of inflation overshooting (i.e., at most 0.5 percentage points above the median) or surging (i.e., at least 2.0 percentage points above the median). The coefficients in Table 5, which quantify how moments change with the labor market, can be combined with the results in panel A of Table 1, which connect moments and inflation probabilities. Using the coefficients in Table 5 from the monthly sample, a one-standard-deviation decline in *ugap* raises the overshoot probability by about a tenth from 21.7 percent to 23.6 percent (to 25.0 percent if using quarterly estimates), while the surge probability declines from 3.4 percent to 2.9 percent (to 2.8 percent with quarterly data). The risk of a surge dips because, as *ugap* compresses, the decline in volatility more than compensates for the increase in kurtosis.

The importance of job market conditions in determining the connection between option-implied inflation moments and the unemployment gap is not driven by the specific choice of how to measure labor tightness. In addition to changes in the participation rate, I consider two alternative proxies: the unemployment gap itself, which implies adding the squared *ugap* to the baseline regressions, and the under-employment rate, defined as the ratio of the under-employment level to the civilian labor force. In both cases, and in contrast to $\Delta part$, lower values correspond to better employment opportunities.

Table 6 shows the coefficients on *ugap* conditional on each of the two measures of labor-market tightness, using monthly data. These results should be juxtaposed to the monthly marginal effects in Table 5. Comparability is ensured because all variables are standardized. There is a remarkable correspondence between the baseline and alternative specifications. Crucially, the coefficients are always larger in tight labor markets. In addition, the only difference in terms of statistical significance pertains to skewness in slack markets, when the proxy is the unemployment gap. In this case, the *t*-statistic is close to -3 , while it is about nil in the baseline. Finally, the magnitude of the coefficients is quite similar to the main results when conditioning on the unemployment gap, while it is noticeably larger if using the under-employment rate.

Table 6. Phillips-Curve Regressions: Alternative Measures of Labor Market Slack

Labor Market	Unemployment Gap		
	vol^{opt}	$skew^{opt}$	$kurt^{opt}$
Slack ($ugap_t = +1\sigma$)	0.180** (2.59)	-0.223*** (-2.82)	0.139 (0.66)
Tight ($ugap_t = -1\sigma$)	0.229** (2.56)	0.503*** (3.86)	-0.301** (-2.18)
Labor Market	Underemployment Rate		
	vol^{opt}	$skew^{opt}$	$kurt^{opt}$
Slack ($underempl_t = +1\sigma$)	0.247 (1.18)	0.521 (1.50)	-0.521 (-1.56)
Tight ($underempl_t = -1\sigma$)	0.339 (0.96)	1.667*** (2.97)	-1.302*** (-2.72)
<p>Note: The table reports marginal coefficients on $ugap_t$ from monthly Phillips-curve regressions similar to those in Table 5. The specifications use two alternative measures of labor market slack: the unemployment gap (meaning that the regressions include a linear and squared term for $ugap_t$) and the underemployment rate ($underempl_t$), defined as the underemployment level divided by the civilian labor force. All variables are standardized. t-statistics are calculated with standard errors based on Newey and West (1987) with 12 lags. The sample covers 2010–19.</p>			

5. The Informative Content of Inflation Moments

The responsiveness of inflation moments to the labor market indicates that options quickly incorporate information, consistent with the discussion in Section 2.2 on why illiquidity does not prevent price formation. Since option-implied distributions are forward looking, moments should anticipate future inflation dynamics, especially—in light of the results so far—when the labor market is tight. In this section, I investigate whether such is indeed the case, using regressions of selected variables on lagged inflation moments. The relatively short sample, which is limited to the post-GFC period, implies that considering the moments of realized inflation as dependent variables is impractical. As a result, I focus on variables that are plausibly linked to lagged inflation moments, as long as option

prices contain useful information. These variables are, first, realized moments of daily risk-neutral inflation expectations and, second, wage-growth dispersion across industries. The following subsections discuss the variables and the results of the analysis.

5.1 *The Moments of Future Break-Even Rates*

I measure risk-neutral inflation expectations with five-year break-even rates, defined as the difference between nominal and real Treasury yields. Realized moments are calculated using daily break-evens over the three months after option-implied moments are measured. Focusing on volatility as an example, if high implied volatility in $t - 1$ correctly anticipates high volatility of realized inflation over the subsequent five years, it is likely that realized inflation will also be volatile over the horizon covered by the break-evens (five years from $t + 2$). As a result, break-evens between t and $t + 2$ are also likely to be volatile. Consequently, high implied volatility in $t - 1$ should anticipate high volatility for break-even rates between t and $t + 2$. Similar reasoning applies to skewness and kurtosis.

I compute the realized log-volatility, skewness, and kurtosis of daily five-year break-even rates (mom_t^{brk}) within a given quarter, half-year, or year. Each of these moments is then regressed on its lagged value and the corresponding lagged average option-implied moment (mom_t^{opt}):

$$mom_t^{brk} = \alpha + \beta \cdot mom_{t-1}^{opt} + \gamma \cdot mom_{t-1}^{brk} + \epsilon_t. \quad (11)$$

In the main specification, the coefficients are computed with ordinary least squares (OLS) and the standard errors are based on Newey and West (1987) with two lags. In an alternative specification, I consider the possibility that persistence in the regressors could yield biased coefficients, and use the procedure of Amihud, Hurvich, and Wang (2009) to correct for this potential bias.

The results are reported in Table 7. Starting with the quarterly horizon, the first three columns of panel A indicate that risk-neutral inflation volatility provides useful signals about future realized break-even volatility, which increases by 0.4 standard deviation for one standard-deviation increase in option-implied volatility. However, there is no effect in the case of skewness or kurtosis. As shown

Table 7. Option-Implied Moments and Future Break-Even Moments

		A. Newey-West											
Method	Horizon	3 Months				6 Months				12 Months			
Realized Moment	→	vol_t^{brk}	$skew_t^{brk}$	$kurt_t^{brk}$	vol_t^{brk}	$skew_t^{brk}$	$kurt_t^{brk}$	vol_t^{brk}	$skew_t^{brk}$	$kurt_t^{brk}$	vol_t^{brk}	$skew_t^{brk}$	$kurt_t^{brk}$
$moment_{t-1}^{opt}$		0.368*** (2.89)	0.083 (0.95)	-0.004 (-0.05)	0.260 (1.58)	0.292*** (2.38)	-0.023 (-0.23)	0.040 (0.22)	0.159 (1.62)	-0.058 (-0.87)			
$moment_{t-1}^{brk}$		0.043 (0.32)	0.091 (1.21)	-0.001 (-0.02)	0.221 (1.56)	-0.103 (-1.26)	-0.005 (-0.05)	0.330 (1.30)	-0.210 (-1.59)	-0.158 (-1.47)			
Obs.		103	103	103	100	100	100	94	94	94			
R ²		0.147	0.015	0.000	0.155	0.094	0.001	0.115	0.075	0.030			
		B. Reduced Bias											
Method	Horizon	3 Months				6 Months				12 Months			
Realized Moment	→	vol_t^{brk}	$skew_t^{brk}$	$kurt_t^{brk}$	vol_t^{brk}	$skew_t^{brk}$	$kurt_t^{brk}$	vol_t^{brk}	$skew_t^{brk}$	$kurt_t^{brk}$	vol_t^{brk}	$skew_t^{brk}$	$kurt_t^{brk}$
$moment_{t-1}^{opt}$		0.391*** (2.91)	0.080 (0.85)	-0.008 (-0.08)	0.273* (1.71)	0.296*** (2.84)	-0.025 (-0.26)	0.049 (0.26)	0.161 (1.38)	-0.058 (-0.86)			
$moment_{t-1}^{brk}$		0.021 (0.16)	0.095 (1.30)	-0.010 (-0.10)	0.210* (1.75)	-0.12 (-1.55)	-0.002 (-0.02)	0.319 (1.15)	-0.217* (-1.85)	-0.158 (-1.61)			

Note: The coefficients reported in the table are from the following regressions using monthly data: $mom_t^{brk} = \alpha + \beta \cdot mom_{t-1}^{opt} + \gamma \cdot mom_{t-1}^{brk} + \epsilon_t$, where mom_t^{brk} is either option-implied log-volatility, skewness, or kurtosis. The variable mom_t^{brk} is either realized log-volatility, skewness, or kurtosis of daily five-year break-even inflation rates over the horizon indicated in the table starting from t . Similarly, the variable mom_{t-1}^{brk} is calculated over the horizon indicated in the table up to and including $t - 1$. In panel A, the coefficients and t -statistics are based on Newey and West (1987) with lags equal to double the horizon. In panel B, the results are based on the methodology of Amihud, Hurvich, and Wang (2009) for persistent regressors. All variables are standardized. The sample covers 2010 to 2019.

in panel B, the reduced-bias methodology of Amihud, Hurvich, and Wang (2009) yields similar coefficients to those based on OLS.

At the half-yearly horizon, the magnitude and statistical significance of the coefficient on volatility start to fade, while risk-neutral skewness is clearly linked to future break-even skewness. Once more, there is little difference between the computations based on OLS and Amihud, Hurvich, and Wang (2009). The connection between break-even and lagged option-implied moments disappears when considering a one-year horizon.

To take the state of the labor market into account, the marginal effect of inflation moments is computed from the coefficients β and λ in the following regression:

$$\begin{aligned} mom_t^{brk} = & \alpha + \beta \cdot mom_{t-1}^{opt} + \gamma \cdot mom_{t-1}^{brk} + \delta \cdot \Delta part_{t-1} \\ & + \lambda \cdot mom_{t-1}^{opt} \cdot \Delta part_{t-1} + \epsilon_t, \end{aligned} \quad (12)$$

where standard errors are based on Newey and West (1987) with two lags. The results are presented in Table 8, which shows marginal effects when the labor market is tight or slack, defined on the basis of $\Delta part_{t-1}$ being one standard deviation above or below average.

At the quarterly horizon, volatility remains the only option-implied moment to comove with its equivalent computed from break-even rates. The picture becomes richer starting with the six-month frequency: labor tightness means that all option-implied moments anticipate their realized counterparts. A similar pattern holds at the one-year horizon, although magnitudes and statistical significance decline for skewness and kurtosis. Interestingly, coefficient signs indicate that inflation risk is persistent in tight labor markets, while it tends to revert in slack ones.

5.2 *The Dispersion of Future Cross-Industry Wage Growth*

I now turn to whether risk-neutral moments anticipate the dispersion of wage growth across industries. The rationale is that inflationary pressures generally raise labor compensation (Blanchard 1986 and Mehra 1991), but industries where workers have low bargaining

Table 8. Option-Implied Moments and Future Break-Even Moments, by Labor Slack

Horizon →	3 Months		6 Months		12 Months	
	Tight	Slack	Tight	Slack	Tight	Slack
vol_{t-1}^{opt}	0.352* (1.81)	0.387 (1.35)	0.400** (2.00)	0.109 (0.38)	0.595*** (3.78)	-0.399** (-2.54)
	<i>Dependent Variable: $vol_{t-1}^{br,k}$</i>					
$skew_{t-1}^{opt}$	0.095 (0.51)	0.080 (0.73)	0.682** (2.29)	0.221** (2.13)	0.665* (1.81)	0.081 (1.13)
	<i>Dependent Variable: $skew_{t-1}^{br,k}$</i>					
$kurt_{t-1}^{opt}$	0.126 (0.67)	-0.042 (-0.37)	0.528* (1.99)	-0.129*** (-2.98)	0.472 (1.09)	-0.141*** (-3.79)
	<i>Dependent Variable: $kurt_{t-1}^{br,k}$</i>					
<p>Note: The coefficients reported in the table are from the following regressions using monthly data: $mom_{t-1}^{br,k} = \alpha + \beta \cdot mom_{t-1}^{opt} + \gamma \cdot mom_{t-1}^{br,k} + \delta \cdot \Delta part_{t-1} + \epsilon_t$, where mom_{t-1}^{opt} is either option-implied log-volatility, skewness, or kurtosis, and $\Delta part_{t-1}$ is the change in the labor force participation rate over the previous 12 months. Tight/slack labor market is defined as $\Delta part_{t-1}$ being one standard deviation above/below the mean. The coefficients and t-statistics are based on Newey and West (1987) with lags equal to double the horizon. All variables are standardized. The sample covers 2010 to 2019.</p>						

power are less likely to see wage growth (Stansbury and Summers 2020). Coupled with downward nominal rigidity in wages, this observation is likely to generate high cross-industry dispersion in wage growth when inflationary pressures are strong.

Dispersion is the log-difference of the highest and lowest wage-growth rates across industries⁶ in a given time period, divided by average growth:

$$disp_t = \ln \frac{\max(w_j) - \min(w_j)}{n^{-1} \cdot \sum_{j=1}^n w_j}, \quad (13)$$

where w_j is the relative change in wages between months t and $t+h$, where $h = 3$ or $h = 6$, depending on the horizon.

To measure worker bargaining power in a given industry, I use the fall in industry employment between 2001 and 2009 (see the discussion of how unemployment affects bargaining power in Summers 1988).⁷ As Pierce and Schott (2016) highlight, employment changes over that period were heavily influenced by a 2000 change in U.S. trade policy that eliminated the risk of higher tariffs on imports from China. Industries most exposed to competition from Chinese imports saw sharp declines in employment.⁸


Figure 2 shows that firms with larger employment declines between 2001 and 2009 experienced lower wage growth between 2010 and 2019, supporting the conjecture that cross-industry dispersion in wage growth depends on labor bargaining power. This finding,

⁶I consider the following industries based on availability (data codes for the St. Louis Fed's FRED data set are shown in parentheses): construction (CES2000000008), education (CES6500000008), financials (CES5500000008), information (CES5000000008), leisure and hospitality (CES7000000008), manufacturing of durable goods (CEU3100000008), manufacturing of nondurable goods (CES3200000008), mining and logging (CES1000000008), professional and business services (CES6000000008), retail trade (CES4200000008), transportation and warehousing (CES4300000008), utilities (CES4422000008), and wholesale trade (CES4142000008).

⁷While union membership is an intuitive measure of bargaining power, union participation in the United States fell from 20 percent in 1983 to 10 percent in 2022 (Bureau of Labor Statistics 2023).

⁸A broad literature finds that import competition normally has adverse effects on wage stability (Bertrand 2004) and distribution (Borjas and Ramey 1995), in line with a negative effect on bargaining power.

Figure 2. Wage Inflation Dispersion and Labor Outside Options



Note: In a given month between 2010 and 2019, industries are sorted into two groups based on wage growth over the following three months. These two groups include, respectively, the three industries with the lowest future wage growth and the industries with the highest future wage growth. As a result, such sorting emphasizes realized cross-industry wage dispersion within time periods. For each group, the bars show the average decline in industry employment between 2001 and 2009, after China was granted permanent normal trade relations by the U.S. Congress in October 2000 (see Pierce and Schott 2016). A smaller decline in employment signals better industry-specific labor outside options during the 2010s.

in turn, suggests that inflationary pressures can affect wage-growth dispersion, since differences in worker bargaining power across industries are reflected in the pass-through from inflationary pressures to wages.

To the extent that implied inflation moments are informative about future inflation risk, they should be positively associated with future wage-growth dispersion, especially if—like high skewness—they express asymmetric risk of high inflation. I gauge the link between risk-neutral inflation moments and future wage-growth dispersion with the following monthly regressions:

$$\begin{aligned} disp_t = & \alpha + \beta \cdot mom_{t-1} + \gamma \cdot disp_{t-1} + \delta \cdot \Delta part_{t-1} \\ & + \lambda \cdot mom_{t-1} \cdot \Delta part_{t-1} + \epsilon_t, \end{aligned} \quad (14)$$

where mom_{t-1} is either log-volatility, skewness, or kurtosis, and standard errors are based on Newey and West (1987) with lags equal to the lead horizon in months plus 12 (the length of the moving window over which $\Delta part$ is calculated). All variables are standardized, so the coefficients can be interpreted as the effect of one-standard-deviation increases in the independent variables on the standard deviation of the dependent variable.

Without conditioning on the labor market, higher volatility and skewness anticipate a wider wage-growth range over the next three months (Table 9). In contrast, kurtosis signals a narrower dispersion. This finding is consistent with the result in Tables 7 and 8 that higher risk-neutral kurtosis, when the coefficient is statistically significant, often prefigures lower realized kurtosis, that is a lower incidence of tail inflation realizations. At the six-month horizon, only the coefficient on skewness remains statistically significant, with a roughly unchanged magnitude. Taking the state of the job market into account (Table 10), the results indicate, once more, that option-implied moments are more informative in periods of labor tightness. The effect is particularly strong for skewness, which, as discussed in the previous paragraph, can be expected to have a stronger effect on wage-growth dispersion because it expresses asymmetric risk of high inflation.

6. Conclusions

While the measured Phillips curve has flattened over time, option-implied inflation distributions indicate that perceived inflation risk has remained sensitive to labor market dynamics. In particular,

Table 9. Inflation Moments and the Dispersion of Industry Wage Growth

Dep. Var. →	<i>disp_t</i>			<i>disp_t</i>		
Horizon →	3 Months			6 Months		
<i>vol_{t-1}^{opt}</i>	0.264** (2.18)			0.060 (0.40)		
<i>skew_{t-1}^{opt}</i>		1.712* (1.75)			1.707* (1.93)	
<i>kurt_{t-1}^{opt}</i>			-0.945** (-2.55)			-0.198 (-0.47)
<i>disp_{t-1}</i>	0.525*** (7.75)	0.554*** (7.88)	0.559*** (8.34)	0.667*** (13.24)	0.637*** (11.12)	0.676*** (13.17)
Obs.	118	118	118	118	118	118
R ²	0.574	0.555	0.564	0.591	0.605	0.590

Note: The dependent variable is the dispersion of wage growth across industries, defined as $disp_t = \ln \frac{\max(w_j) - \min(w_j)}{n-1 \cdot \sum_{j=1}^n w_j}$, where w_j is the relative change in wages between months t and $t + 3$ or $t + 6$, depending on the horizon indicated in the panels above. The variable $disp_{t-1}$ is defined in a similar way, but with the relative change in wages being computed between months $t - 3$ or $t - 6$ and t , depending on the horizon. All variables are standardized and t -statistics are computed using standard errors based on Newey and West (1987) with lags equal to twice the lead/lag horizon. The monthly sample covers 2010–19.

Table 10. Inflation Moments and Wage-Growth Dispersion, by Labor Slack

Horizon →	Marginal Effect of Option-Implied Moments			
	3 Months		6 Months	
Labor Market →	Tight	Slack	Tight	Slack
<i>vol_{t-1}^{opt}</i>	0.303* (1.88)	0.089 (0.47)	0.061 (0.64)	-0.013 (-0.05)
<i>skew_{t-1}^{opt}</i>	0.494*** (4.98)	0.062 (1.29)	0.228** (2.21)	0.118** (2.05)
<i>kurt_{t-1}^{opt}</i>	-0.306** (-2.20)	-0.151** (-2.11)	0.067 (0.57)	-0.087 (-1.17)

Note: This table reports the coefficients of regressions (monthly frequency) linking the future dispersion of industry wage growth to lagged option-implied inflation moments. The specification is $disp_t = \alpha + \beta disp_{t-1} + \gamma x_{t-1} + \delta \cdot \Delta part_{t-1} + \lambda \cdot x_{t-1} \cdot \Delta part_{t-1} + \epsilon_t$. The variable x_{t-1} represents one of the different option-implied moments, while $disp_t$ and $disp_{t-1}$ are defined as detailed in Table 9. All variables are standardized and t -statistics are computed using standard errors based on Newey and West (1987) with lags equal to the lead horizon plus 12 ($\Delta part$ is computed over 12 months). The sample covers 2010–19.

upside risk increases when the unemployment gap shrinks, especially in an already tight labor market. Among the various inflation moments, implied skewness is particularly reactive to the unemployment gap. Importantly, implied moments convey useful information, in that they anticipate future patterns in market-based inflation expectations and in wage growth. Overall, the results presented in this paper point to complex dynamics for (risk-neutral) inflation risk in the background of an apparent detachment of headline inflation from the labor market.

The use of options to compute inflation moments has two important implications. The first is that the moments incorporate risk premia, which generally create a wedge relative to estimates based on surveys or econometric methods. These premia, however, incorporate useful information on outcomes that matter the most to investors and that affect their decision-making. The second consequence is that the reliability of the implied moments depends on the quality of option prices. While these instruments were illiquid in the second half of the sample and prices largely reflected dealer quotes, dealers have strong incentives to provide quotes close to fair value. The main reason is that their clients reportedly use the data to mark-to-market legacy positions in inflation options.

Appendix

A.1 Effects of Including Data from the Pandemic

This section provides more details on the behavior of inflation moments during the pandemic and on the effect of including the pandemic in the sample on the main results. While data are available to cover the pandemic and its immediate aftermath, the paper focuses on the pre-pandemic period for two main reasons.

The first reason is that the dislocations generated by the pandemic and the resulting data outliers can distort inference about economic relations that hold during normal periods. There has been considerable effort in the literature to contain these distortions. See, for instance, Davis and Ng (2023) and Carriero et al. (2024). The second reason is that, as noted in the main text, institutional features of the market for inflation options mean that the informativeness of these options was expected to decline through the early 2020s.

Table A.1. Marginal Effects of Unemployment Gap

Labor Market	Marginal Effects		
	Vol.	Skew	Kurtosis
Slack	0.170 (0.97)	-0.547* (-1.75)	0.091 (0.30)
Tight	0.376*** (2.88)	0.936*** (2.76)	-0.615** (-2.43)

Indeed, data for several strikes were discontinued starting in August 2021.

The remainder of this section provides selected details on how including data from the pandemic (up to August 2021) affects the analysis. It will then provide details on the liquidity of inflation options in the latter part of the sample.

Table A.1 shows the marginal effects of the unemployment gap on the various implied moments when the labor market is slack or tight, when using quarterly data that extends to the end of the sample in the second quarter of 2021. The table corresponds to the bottom left panel of Table 5 in the manuscript. The results are quite similar to those shown in Table 5.

While changes in data availability in August 2021 do not affect the analysis discussed in the paper, which focuses on the pre-pandemic sample, a pertinent question is whether the quality of the data started to decline before the quotes for some strikes became unavailable. I address this point in the first part of the next section.

A.2 Evidence on Options Data Informativeness

A.2.1 The Effect of Discarding Stale Prices

Even if no volume is available for inflation options, it is possible to gauge the quality of the data by studying instances when prices do not change from one day to the other (“stale” prices). Prices that remain unchanged are not necessarily a sign of poor data quality. Especially for out-of-the-money (OTM) options, which by construction have low prices, a stale quoted price can simply mean that the

intrinsic value of the option has not changed more than the minimum tick used by an exchange or by a desk that disseminates quotes.

Before proceeding further, it is useful to consider the incidence of prices that do not change from one day to the next in the market for options on the S&P 500 index (ticker SPX), which are some of the most liquid options available. The following statistics refer to data for the month of August 2023, obtained from Optionmetrics via Wharton Research Data Services.

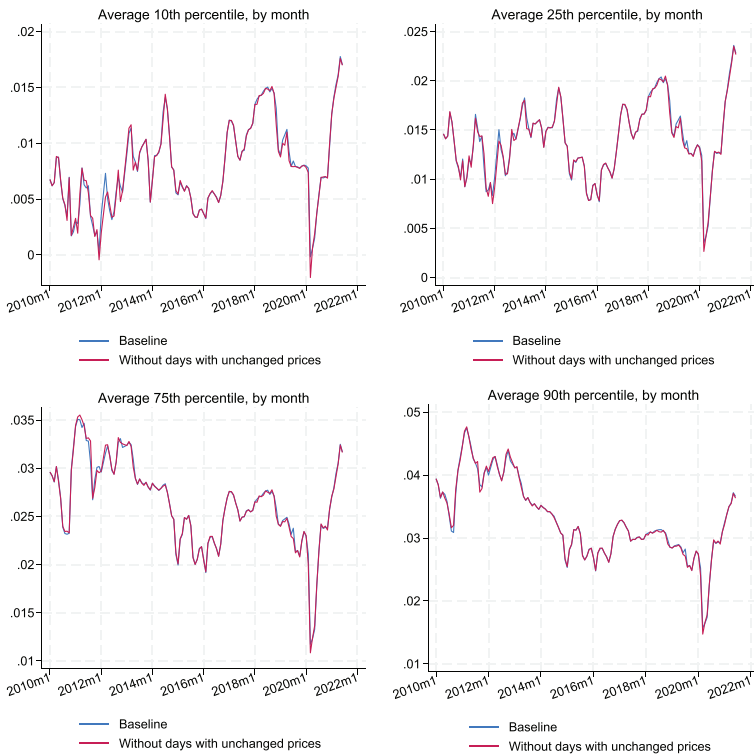
- There are 400,958 observations in the August 2023 file.
- Of these, 29,211 observations (7 percent) have either the bid price or ask price that did not change compared to the previous trading day.
- Of these, 28,289 observations (97 percent) are OTM options (with the absolute value of the option delta below 0.1). The high incidence of OTM options is expected, because low delta options have low prices and the “true price” may move less than the minimum tick.
- Of these 28,289 observations, 7,204 (25 percent) have non-zero trading volume, and 4,325 (15 percent) have trading volume greater than 10 contracts

Going back to the question of whether the quality of the inflation-option data started to decline before the quotes for some strikes became unavailable, a conservative approach to gauging the issue is to discard days in which at least one of the options has a price that does not change relative to the previous day. Figure A.1 shows option-implied inflation percentiles computed this way (red lines) along with the percentiles based on the data used in the manuscript (blue lines). Besides a discrepancy in the 10th percentile at the beginning of 2012, the baseline moments are virtually identical to their counterparts that drop days when at least one option price is unchanged. This is also the case before and during the pandemic.

A.2.2 Evidence from Option-Implied Volatility

The implied volatility of inflation saw large movements around the pandemic (see Figure 1 in the main text). While this behavior is


Figure A.1. Implied Inflation Percentiles after Removing Stale Prices



not unexpected given the impact of the pandemic, it is important to ensure that these movements are not the consequence of issues with the data or methodology. Even if the pandemic is not included in the sample and cannot drive the results, assessing the dynamics of implied moments around the pandemic is useful to confirm that the data and methodology are sound.

Figure A.2 shows a scatterplot of implied inflation volatility (x-axis) against a measure of expected-inflation dispersion obtained from the Survey of Professional Forecasters (SPF), provided by the Federal Reserve Bank of Philadelphia (y-axis). The measure of dispersion is the interquartile range of forecasts for inflation over the next five years. The SPF data are quarterly, and implied inflation

Figure A.2. Inflation Survey Dispersion and Implied Inflation Volatility



volatility is averaged within a quarter. The comovement is relatively strong between 2010 and 2019 (the main sample) and the correlation is 68 percent. Including data through the second quarter of 2021, the correlation drops to 63 percent, but the comovement is still quite strong. The one exception is the first quarter of 2021, when option-implied volatility is much higher than the survey-based measure.

Figure A.3 depicts the time series of daily implied inflation volatility during the first quarter of 2020, when asset prices were extremely volatile due to the initial impact of the pandemic. Early in the quarter, implied volatility was in line with end-2019 values, but rose to very high levels in mid-March 2020 and declined only after the Federal Reserve set up several liquidity facilities in support of credit intermediation. This spike is the reason why the average value of inflation implied volatility in the first quarter of 2020 is unusually high. Other measures of financial market risk jumped to historic highs around that time. For instance, the VIX index closed at 208 percent on March 16, 2020, about 15 times the average closing value in the fourth quarter of 2019 (14 percent).

Figure A.3. Option-Implied Volatility during the First Quarter of 2020



Note: The shaded area covers March 17 to March 23, 2020. During this period, the Federal Reserve announced several liquidity facilities to support credit intermediation.


A.2.3 Implied Distribution Dynamics in Early 2020

Reviewing inflation densities during the pandemic can help to better understand whether information embedded in option prices reflects rapid shifts in economic conditions during a challenging period. Figure A.4 shows the time series of percentiles of implied densities, which make it easier to understand changes in the distributions over time.

The chart includes selected percentiles ranging from the 10th to the 90th, spanning 2020. The chart also shows March 23, 2020, when the Federal Reserve set up several liquidity facilities that supported credit markets. Starting in mid-February, implied distributions start shifting rapidly to the left as the impact of the pandemic had large negative effects on current and expected economic growth. The shift to the left was precipitous in March and started to revert as policymakers deployed significant resources to contain the economic fallout of the pandemic. By late summer, the distributions returned to pre-pandemic shapes.

Figure A.5 shows the overall change in percentiles between January 2 and March 23 (corresponding to the vertical line shown in the time series in Figure A.4). The bar chart represents the shift in the shape of the implied distribution. While the whole distribution

Figure A.4. Selected Percentiles of Option-Implied Inflation Distributions in 2020



Note: The vertical line shows March 23, 2020, when the Federal Reserve established liquidity facilities in support of corporate credit intermediation.


shifted left, as indicated by the drop in the median, the right tail compressed even further, as shown by the large decline in the 90th percentile.

A.2.4 Implied Moments after Key Monetary Policy Announcements

This section provides evidence on how implied inflation moments react to important announcements that shifted the monetary policy stance. Specifically, I compute the average change in implied volatility, skewness, and kurtosis following certain announcements that changed the Federal Reserve's balance sheet policy.⁹ I then compare these averages to average changes over the rest of the sample. As shown in Figure A.6, average month-on-month changes in implied moments after these events (blue dots) are outside of the 90

⁹The events are QE2 hinted (August 2010), Operation Twist announced (September 2011), Operation Twist extended (June 2012), QE3 expanded to include more Treasuries (December 2012), QE3 tapering begins (December 2013), QE3 ends (October 2014), Federal Reserve announces balance sheet normalization (June 2017), and Federal Reserve Chair emphasizes balance sheet flexibility if needed (January 2019).

Figure A.5. Change in Implied Percentiles, from January 2 to March 23, 2020



Note: On March 23, 2020, the Federal Reserve established liquidity facilities in support of corporate credit intermediation.


percent confidence intervals for average changes in the rest of the sample, except for kurtosis (top chart). While average changes over the full sample can be expected to be close to zero (the moments should be stationary), the same conclusions apply when considering the absolute values of changes (bottom chart).

On average, volatility decreases and skewness increases, meaning that—following significant monetary policy announcements—investors become less uncertain about the future path of inflation, and (due to the increase in skewness) the risk of low inflation subsides.

A.2.5 Implied Distributions after Selected QE Announcements

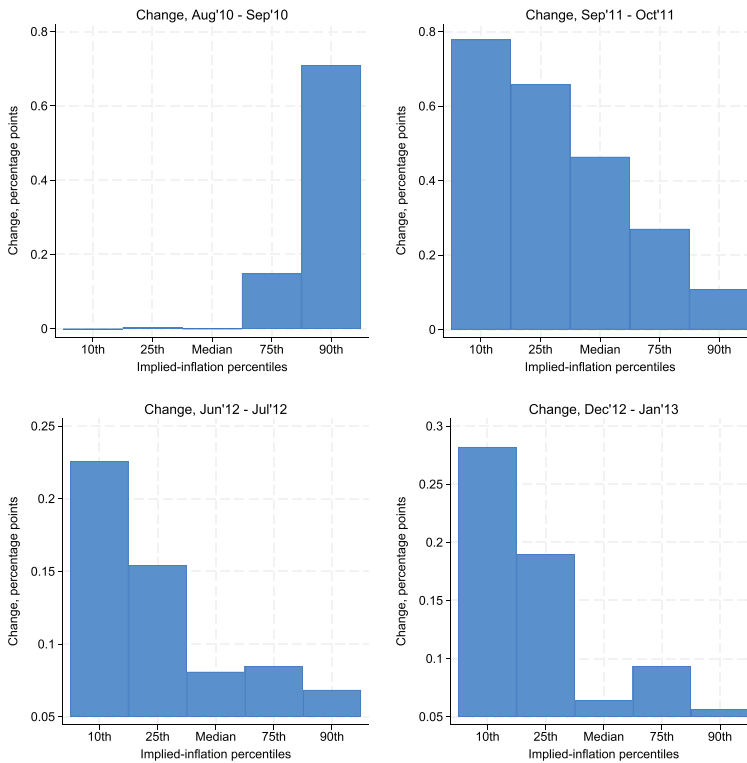
To illustrate the effect that significant announcements related to the monetary policy stance have on the implied distribution of inflation, the four charts in Figure A.7 show changes in selected percentiles

Figure A.6. Changes in Implied Moments after Key Monetary Policy Announcements



from the month in which an announcement was made to the following month. The charts refer to the following announcements: QE2 hinted (August 2010); Operation Twist announced (September 2011); Operation Twist extended (June 2012); and QE3 expanded to include more Treasuries (December 2012). In the first case, higher percentiles increased markedly, while the rest remained unchanged. As such, the distribution saw a thickening of the right tail. In contrast, each of the three other announcements saw an increase in the

Figure A.7. Changes in Implied Distribution Percentiles after QE Announcements



lower percentiles, indicating a shift in the left-tail probability mass toward the median and a declining risk of low inflation.

References

- Aït-Sahalia, Y., and J. Duarte. 2003. "Nonparametric Option Pricing under Shape Restrictions." *Journal of Econometrics* 116 (1–2): 9–47.
- Amihud, Y., C. Hurvich, and Y. Wang. 2009. "Multiple-Predictor Regressions: Hypothesis Testing." *Review of Financial Studies* 22 (1): 413–34.
- Andrade, P., E. Ghysels, and J. Idier. 2015. "Tails of Inflation Forecasts and Tales of Monetary Policy." Working Paper.

- 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.
- Banerjee, R., J. Contreras, A. Mehrotra, and F. Zampolli. 2020. "Inflation at Risk in Advanced and Emerging Market Economies." BIS Working Paper No. 883.
- Barnichon, R., and G. Mesters. 2021. "The Phillips Multiplier." *Journal of Monetary Economics* 117 (January): 689–705.
- Bertrand, M. 2004. "From the Invisible Handshake to the Invisible Hand? How Import Competition Changes the Employment Relationship." *Journal of Labor Economics* 22 (4): 723–65.
- Blanchard, O. 1986. "The Wage Price Spiral." *Quarterly Journal of Economics* 101 (3): 543–66.
- . 2016. "The Phillips Curve: Back to the '60s?" *AEA Papers and Proceedings* 106 (5): 31–34.
- Borjas, G., and V. Ramey. 1995. "Foreign Competition, Market Power, and Wage Inequality." *Quarterly Journal of Economics* 110 (4): 1075–1110.
- Bowley, A. 1920. *Elements of Statistics*. New York: Charles Scribner's Sons.
- Breeden, D., and R. Litzenberger. 1978. "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of Business* 51 (4): 621–51.
- Bureau of Labor Statistics. 2023. "Union Membership Rate Fell by 0.2 Percentage Point to 10.1 Percent in 2022." *The Economics Daily*, January 24.
- Busetti, F., M. Caivano, and L. Rodano. 2015. "On the Conditional Distribution of Euro Area Inflation Forecasts." Working Paper.
- Carriero, A., T. E. Clark, M. Marcellino, and E. Mertens. 2024. "Addressing COVID-19 Outliers in BVARs with Stochastic Volatility." *Review of Economics and Statistics* 106 (5): 1403–17.
- 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): 197–232.

- Coibion, O., Y. Gorodnichenko, and R. Kamdar. 2018. "The Formation of Expectations, Inflation, and the Phillips Curve." *Journal of Economic Literature* 56 (4): 1447–91.
- Davis, R., and S. Ng. 2023. "Time Series Estimation of the Dynamic Effects of Disaster-Type Shocks." *Journal of Econometrics* 235 (1): 180–201.
- Erceg, C., and A. Levin. 2014. "Labor Force Participation and Monetary Policy in the Wake of the Great Recession." *Journal of Money, Credit and Banking* 46 (S2): 3–49.
- Eser, F., P. Karadi, P. Lane, L. Moretti, and C. Osbat. 2020. "The Phillips Curve at the ECB." *Manchester School* 88 (S1): 50–85.
- Feldman, R., K. Heinecke, N. Kocherlakota, S. Schulhofer-Wohl, and T. Tallarini. 2015. "Market-Based Probabilities: A Tool for Policymakers." Report, Federal Reserve Bank of Minneapolis.
- Fleckenstein, M., F. Longstaff, and H. Lustig. 2017. "Deflation Risk." *Review of Financial Studies* 30 (8): 2719–60.
- Galati, G., Z. Gorgi, and R. M. Zhou. 2018. "Deflation Risk in the Euro Area and Central Bank Credibility." *Economics Letters* 167 (June): 124–26.
- Gizatulina, A., and Z. Hellman. 2019. "No Trade and Yes Trade Theorems for Heterogeneous Priors." *Journal of Economic Theory* 182 (July): 161–84.
- Hilscher, J., A. Raviv, and R. Reis. 2022. "Inflating Away Public Debt? An Empirical Assessment." *Review of Financial Studies* 35 (3): 1553–95.
- . 2026. "How Likely is an Inflation Disaster?" *Review of Financial Studies* 39 (3): 744–82.
- Hobijn, B., and A. Şahin. 2021. "Maximum Employment and the Participation Cycle." In *Macroeconomic Policy in an Uneven Economy*. Proceedings of the 2021 Economic Policy Symposium. Kansas City, MO: Federal Reserve Bank of Kansas City.
- Hooper, P., F. 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.
- Juselius, M., and E. Takáts. 2021. "Inflation and Demography through Time." *Journal of Economic Dynamics and Control* 128 (July): Article 104136.

- Kitsul, Y., and J. Wright. 2013. "The Economics of Options-Implied Inflation Probability Density Functions." *Journal of Financial Economics* 110 (3): 696–711.
- Krasker, W. 1980. "The 'Peso Problem' in Testing the Efficiency of Forward Exchange Markets." *Journal of Monetary Economics* 6 (2): 269–76.
- Li, G. 1985. "Robust Regression." In *Exploring Data Tables, Trends, and Shapes*, ed. D. C. Hoaglin, F. Mosteller, and J. W. Tukey. New York: Wiley.
- Lombardi, M., M. Riggi, and E. Viviano. 2020. "Bargaining Power and the Phillips Curve: A Micro-Macro Analysis." BIS Working Paper No. 902.
- López-Salido, D., and F. Loria. 2024. "Inflation at Risk." *Journal of Monetary Economics* 145 (Supplement): Article 103570.
- Manzan, S., and D. Zerom. 2013. "Are Macroeconomic Variables Useful for Forecasting the Distribution of U.S. Inflation?" *International Journal of Forecasting* 29 (3): 469–78.
- McLeay, M., and S. Tenreyro. 2020. "Optimal Inflation and the Identification of the Phillips Curve." *NBER Macroeconomics Annual 2019*, Vol. 34, ed. M. S. Eichenbaum, E. Hurst, and J. A. Parker, 199–255. University of Chicago Press.
- Mehra, Y. 1991. "Wage Growth and the Inflation Process: An Empirical Note." *American Economic Review* 81 (4): 931–37.
- Mercurio, F., and J. Zhang. 2017. "Pricing Inflation Derivatives." Bloomberg LP.
- Mertens, T., and J. Williams. 2021. "What to Expect from the Lower Bound on Interest Rates: Evidence from Derivatives Prices." *American Economic Review* 111 (8): 2473–2505.
- Milgrom, P., and N. Stokey. 1982. "Information, Trade and Common Knowledge." *Journal of Economic Theory* 26 (1): 17–27.
- Mojon, B., and X. Ragot. 2019. "Can an Ageing Workforce Explain Low Inflation?" BIS Working Paper No. 776.
- Moors, J. 1988. "A Quantile Alternative for Kurtosis." *The Statistician* 37 (1): 25–32.
- Nagel, S. 2016. "Long-Run Inflation Uncertainty." *International Journal of Central Banking* 12 (3): 207–17.
- Newey, W., and K. West. 1987. "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–8.

- Pierce, J., and P. Schott. 2016. "The Surprisingly Swift Decline of US Manufacturing Employment." *American Economic Review* 106 (7): 1632–62.
- Ratner, D., and J. Sim. 2020. "Who Killed the Phillips Curve? A Murder Mystery." Working Paper.
- Reis, R. 2020. "The People versus the Markets: A Parsimonious Model of Inflation Expectations." Working Paper.
- Santa-Clara, P., and S. Yan. 2010. "Crashes, Volatility, and the Equity Premium: Lessons from S&P 500 Options." *Review of Economics and Statistics* 92 (2): 435–51.
- Stansbury, A., and L. Summers. 2020. "The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy." *Brookings Papers on Economic Activity* (Spring): 1–77.
- Summers, L. 1988. "Relative Wages, Efficiency Wages, and Keynesian Unemployment." *AEA Papers and Proceedings* 78 (2): 383–88.
- Tobin, J. 1972. "Inflation and Unemployment." *American Economic Review* 62 (112): 1–18.
- Watson, M. 2014. "Inflation Persistence, the NAIRU, and the Great Recession." *AEA Papers and Proceedings* 104 (5): 31–36.
- Wolfers, J., and E. Zitzewitz. 2004. "Prediction Markets." *Journal of Economic Perspectives* 18 (2): 107–26.
- Yellen, J. 2014. "Labor Market Dynamics and Monetary Policy." Speech at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, August 22.

Monetary Policy Shocks and the Employment of Young, Middle-Aged, and Old Workers*

Fumitaka Nakamura, Nao Sudo, and Yu Sugisaki
Bank of Japan

We study how monetary policy affects the labor status of people of different ages and genders, using Japanese data spanning the late 1990s to the late 2010s, with monetary policy shocks identified using high-frequency market data. We first demonstrate that expansionary monetary policy shocks reduce the number of unemployed of all ages in both genders by almost the same amount. We then show that the impacts of these shocks are starkly different across ages in terms of responses in the labor force and number of employed. Specifically, expansionary monetary policy shocks induce the young and elderly demographics who were previously outside the labor force to enter the labor market, thereby increasing the employed population within these age brackets while exerting lesser influence on the middle-aged cohort. These findings suggest that changes in the labor force participation rate could play an important role in determining the degree of labor market slack for specific ages, potentially leading to a relatively limited wage increase to expansionary monetary policy shocks through the composition effect.


JEL Codes: E24, E32, E52.

1. Introduction

The demographic landscape of the Japanese economy has undergone significant transformations. Figure 1 shows the population

*We would like to thank Daisuke Miyakawa, Makoto Nirei, Toshitaka Sekine, Mototsugu Shintani, seminar participants at the Macroeconomic Workshop at Osaka University, and the staff of the Bank of Japan for their valuable comments and suggestions. The views expressed in this paper are those of the authors and do not necessarily represent the views of the Bank of Japan. Author e-mails: fumitaka.nakamura@boj.or.jp; nao.sudou@boj.or.jp; yuu.sugisaki@boj.or.jp.

Figure 1. Distribution of Population by Age in Japan



Note: This figure shows the distribution of population by age in Japan. The solid line, the dotted line, and the dashed line show the distribution in 1970, 1995, and 2019, respectively. The data come from the Ministry of Internal Affairs and Communications.

distribution in Japan across pivotal years, 1970, 1995, and 2019. In 1970, the distribution peaked around the 20s, with individuals younger than 30 constituting more than half of the population. In contrast, by 2019, this demographic segment accounted for less than 30 percent, while individuals aged over 60 represented one-third of the total population. Indeed, Japan is not an exception; to some degree, other advanced economies such as the United States, the United Kingdom, and Germany have also encountered analogous changes in their demographic structures in the last 50 years to comparable extents.

The effectiveness of macroeconomic policies is intricately linked to the demographic landscape, particularly in the context of labor market dynamics where individuals transition in and out of the workforce. Responses in labor force participation to policy changes across different stages of the life cycle significantly affect policy outcomes (see İmrohoroğlu and Kitao 2012 and Leahy and Thapar 2022). For example, younger demographics contemplating entry into the workforce and elderly individuals pondering retirement decisions are influenced by economic, social, and institutional factors. Monetary policy shocks can also expedite or delay these decisions by altering economic conditions, leading to heterogeneous responses of the individuals across age groups through affecting the factors involved in


their decision-making process: some age cohorts may possess greater flexibility in adjusting their labor status compared to the other age cohorts. Because of the differences regarding the responses across age groups, the demographic composition can crucially affect monetary policy outcomes regarding total labor inputs, wage growth, and aggregate output. As illustrated in Figure 2, in Japan, younger and elderly individuals typically have earned lower wages compared to their middle-aged counterparts in the last 20 years, although the differences become slightly diminished in 2019 compared to 1999.¹ Consequently, other things being equal, an expansionary monetary policy shock may result in a limited nominal wage increase if it predominantly increases low-wage workers such as the young and the elderly relative to middle-aged workers. Therefore, it may also result in a moderate response of inflation as long as institutional arrangements in Japan that make both young and old workers responsive to these shocks remain in place.²

First, this paper investigates the response to monetary policy shocks of individuals with different ages concerning labor status.³ To address this question, we utilize disaggregated data on employment, unemployment, and labor force participation by age and gender obtained from the Labor Force Survey in Japan, a country characterized by having one of the highest proportions of elderly citizens globally. We employ a vector autoregression (VAR) framework to estimate the responses of these disaggregated labor-related variables to monetary policy shocks. Our empirical specification closely follows the approach developed by Gertler and Karadi (2015) and identifies monetary policy shocks by instrumenting the one-year government

¹ Growing participation of the younger and elderly workers in the labor force may have suppressed wage growth during the recovery from the Global Financial Crisis. For example, for elderly workers, Mojon and Ragot (2019) and Bank of Japan (2023) highlight that active participation of old workers boosts the labor force despite the overall population declines, which contributes to a moderate increase in wages driven by the increased labor supply in the 2010s for Japan and OECD countries, respectively.

²This is consistent with Imam (2015), which shows the weakening effect of monetary policy shocks on the inflation rate due to aging.

³The importance of demographic change in monetary policy decision-making is also pointed out, for example, in Geneakoplos, Magill, and Quinzii (2004), Yellen (2014), Favero, Gozluklu, and Yang (2016), and Favero, Melone, and Tamoni (2022).

Figure 2. Hourly Nominal Earnings by Age in Japan

Note: This figure shows the average hourly nominal earnings by age in Japan. The data come from the Ministry of Health, Labour and Welfare.

bond yield with changes in Euroyen futures (Japanese short-term interest rate futures), following the methodology outlined by Nakamura, Sudo, and Sugisaki (2024).

We find that the number of unemployed falls significantly in response to an expansionary monetary policy shock, and the magnitude of the decline is quantitatively the same across ages and genders. In terms of changes in the labor force and number of employed, however, the impacts of a monetary policy shock are starkly different across ages. Specifically, both youth and elderly individuals experience a greater increase in labor force participation compared to the middle-aged demographic in response to expansionary monetary policy shocks, suggesting that a favorable macroeconomic labor market environment lowers the relative benefits of being a student or a retiree for individuals of these age cohorts. Moreover, the number of employed increases significantly for the youth and the elderly while that of the middle-aged does not respond significantly to the shocks. This observation suggests that a rise in firms' increased demand for labor inputs following expansionary monetary policy shocks is met, if not fully, by a rise in participation in the labor force among young and elderly individuals.

We further delve into the distinctive characteristics of Japan's labor market structure to illuminate the factors behind the pronounced increases in both young and older workers' responses to

expansionary monetary policy shocks. Regarding young individuals, their labor force participation tends to be procyclical as they transition between students and workers; economic upswings encourage a greater number of students to enter the labor market, while recessions prompt a withdrawal. This procyclicality is also partially attributed to the fact that schooling options help youths choose better timing for employment. Under a system of simultaneous hiring and lifetime employment, young people have an incentive to be selective about the first company they work for as new graduates and the type of work they do, such as permanent or temporary employment from a long-term perspective. Regarding the elderly population, their labor force participation tends to be procyclical as they transition between retirees and workers. With a significant perception of inadequate pension support, coupled with high life expectancy and financial insecurity, the retired elderly rejoin the labor force when the labor market condition is favorable.

Second, we estimate the wage Phillips curve for different age groups to evaluate how the fluctuations in the labor force participation rate affect the labor market slack by altering the available labor supply within the economy. When the cyclicity of changes in the rate is age-specific, as evidenced in our study, variations in the slope of the wage Phillips curve are expected to be different across the age groups. Indeed, we find that the slope is steeper for the middle-aged compared to the slopes for the youth and the elderly. This observation further suggests that the aging of society has gone hand-in-hand with weak wage dynamics.⁴ Our findings are consistent with the existing studies, for example, by Erceg and Levin (2014) and Elsby, Hobijn, and Şahin (2015): the shifts in labor force participation rate play a pivotal role in shaping nominal wage dynamics. Indeed, Erceg and Levin (2014), stressing the role of the labor force participation rate, showed theoretically that the slope of the wage Phillips curve becomes less steep when the labor force participation rate exhibits cyclical patterns.

⁴The flattening of the wage Phillips curve has been reported in the existing works in advanced countries, including Japan. See, for example, an empirical study conducted by Iwasaki, Muto, and Shintani (2021) and Hirata, Maruyama, and Mineyama (2020).

This paper is organized as follows. Section 2 describes the related literature. Section 3 presents the econometric framework we use for the empirical analysis. Section 4 describes the data set. Section 5 documents the results. Section 6 discusses the implications of our results for the wage Phillips curve. Section 7 concludes.

2. Literature Review

This paper is related to three strands of the literature. The first strand covers empirical works quantifying the heterogeneous effects of monetary policy shocks depending on the characteristics of each agent, including studies such as Coibion et al. (2017), Inui, Sudo, and Yamada (2017), Bahaj et al. (2019), Cloyne, Ferreira, and Surico (2020), Zens, Böck, and Zörner (2020), and Leahy and Thapar (2022). Among these studies, Bahaj et al. (2019) and Zens, Böck, and Zörner (2020) investigate the relationship between monetary policy shocks and the response to employment. Bahaj et al. (2019) focus on the collateral channel by exploiting the homes of firms' directors as a key source of collateral. The spatial separation of firms from their collateral enables them to separate the collateral channel from local demand effects. Zens, Böck, and Zörner (2020) find that the effects of monetary policy shocks on the unemployment rate differ depending on the occupational group. Compared to these papers, which focus on the characteristics of companies, this paper focuses on the heterogeneous effects of monetary policy shocks depending on differences in workers' age and gender. Moreover, Coibion et al. (2017) and Inui, Sudo, and Yamada (2017) investigate the relationship between consumption inequality and the transmission of monetary policy in the U.S. and in Japan, respectively. Inui, Sudo, and Yamada (2017) find that monetary policy shocks do not have statistically significant impacts across Japanese households in a stable manner in terms of their effects on income and consumption inequality. On the other hand, the present paper discusses the relationship between monetary policy and labor status and finds that impacts are starkly different across ages.⁵

⁵Regarding the relationship between monetary policy and inequality, Bernanke (2015) points out the possibility that quantitative easing can reduce inequality through promoting job creation since a stronger labor market benefits

The second strand of the literature covers empirical work quantifying the relationship between population aging and the business cycle, including the implications for monetary policy, such as Jaimovich and Siu (2009), Imam (2015), and Wong (2019).⁶ Among these studies, Jaimovich and Siu (2009) document the relatively larger variation of employment among the youth and the elderly, and argue that the change in the age composition of the labor force accounts for a significant fraction of the cyclical fluctuation of business cycles in the U.S. We show that a similar pattern can be seen for employment variations conditional on monetary policy shocks in Japan as well as in the U.S., namely a larger response of the number of employed among the young and the old. Imam (2015) shows that there is a relationship between population aging and monetary policy in terms of inflation and unemployment by Bayesian estimation techniques, while this paper assesses the effect of monetary policy shocks on labor-related variables of different age groups.

The third strand of the literature consists of studies investigating the labor market slack and wage settings, including studies such as Erceg and Levin (2014), Blanchflower and Levin (2015), Elsby, Hobijn, and Şahin (2015), and Christiano, Trabandt, and Walentin (2021).⁷ Erceg and Levin (2014) argue theoretically that monetary policy affects the labor force participation rate. Christiano, Trabandt, and Walentin (2021) construct a monetary model in which the labor force participation rate varies with the business cycle. Blanchflower and Levin (2015) and Elsby, Hobijn, and Şahin (2015) empirically point out the importance of labor force participation margin in explaining the hidden unemployment, which could put downward pressure on nominal wages. Since some workers who are not actively searching for a job can rejoin the workforce if the job

the middle class. This paper clarifies the segments of the population who receive benefits from expansionary monetary policy shocks: such shocks lead to a greater increase in employment of young and old workers, whose wages are relatively low.

⁶Instead of monetary policy, some studies investigated the relationship between fiscal policy and demographic change, since age structure is closely related to the social security policy and taxation system. This includes Ríos-Rull (2001), Abel (2003), and İmrohoroğlu and Kitao (2012).

⁷Also, a lot of study has been dedicated to clarifying the cyclicity of employment and labor force participation. See, for instance, Ragan (1977), Shimer (2005), Campolmi and Gnocchi (2016), Chodorow-Reich and Karabarbounis (2016), and Cairo, Fugita, and Morales-Jimenez (2022).

market gets stronger, labor market slack exerts significant downward pressure on nominal wages. Compared to these studies, we argue that the slack is age-specific, larger for the young and the old, and smaller for other age groups, which, together with the shift in age composition in the population, may moderate a rise in the aggregate-level nominal wage in Japan in response to expansionary monetary policy shocks in Japan.

3. Econometric Framework

Our estimation procedure consists of two steps. First, we formulate a VAR that consists of macroeconomic variables of the sample period running from 1999:M1 to 2018:M12 and estimate the response of the variables to monetary policy shocks. We identify monetary policy shocks by using high-frequency data as external instruments following Gertler and Karadi (2015). Almost all of our sample period coincides with the period when the short-term interest rate was set close to zero and the Bank of Japan had launched various policy initiatives trying to reduce not only the current short-term nominal interest rate but also the expected short-term nominal interest rate and the term premium.⁸ We therefore incorporate financial variables including the government bond rate instead of the short-term nominal interest rate in our VAR and isolate variations in the government bond rate that stem from the exogenous policy shocks thanks to high-frequency identification.⁹ Second, we regress a disaggregated labor-related variable, such as the unemployment rate of males in their 30s, on macroeconomic variables and estimate the response of the variable of interest to monetary policy shocks. We assume a block

⁸Our data set includes the period in which the zero interest rate (ZIR), quantitative easing (QE), comprehensive monetary easing (CME), and quantitative and qualitative monetary easing (QQE) policies were in place. Note also that QQE includes QQE with a negative interest rate and QQE with yield curve control. ZIR runs from 1999:M2 to 2000:M8, QE runs from 2001:M3 to 2006:M3, CME runs from 2010:M10 to 2013:M4, and QQE runs from 2013:M4 to the present.

⁹External instruments are employed in a growing body of studies, including Stock and Watson (2012) and Mertens and Ravn (2013). High-frequency identification for Japan's monetary policy shocks used in the current paper are the same as those constructed in Nakamura, Sudo, and Sugisaki (2024).

recursive framework (see Lee and Ni 2002) so that a shock to disaggregated variables does not affect the dynamics of macrovariables in the VAR estimated in the first step.

3.1 Macrovariable Block

Our VAR contains N_x number of monthly series of macroeconomic variables and is expressed in the following structural form:

$$A_x X_t = A_x c_x + \sum_{j=1}^p A_x B_{xx,j} X_{t-j} + \varepsilon_{x,t}, \quad (1)$$

where X_t is the N_x -dimensional vector of macroeconomic variables, $B_{xx,j}$ is the matrix of polynomials that represent the structural relationships among macroeconomic variables, A_x and c_x are coefficient matrices, and $\varepsilon_{x,t}$ is a vector of structural white noise shocks.¹⁰ Multiplying each side of the equation by A_x^{-1} yields the reduced-form VAR model given by

$$X_t = c_x + \sum_{j=1}^p B_{xx,j} X_{t-j} + u_{x,t}, \quad (2)$$

where $u_{x,t}$ is the reduced-form shock. The relationship between the reduced-form shock, $u_{x,t}$, and the structural shock, $\varepsilon_{x,t}$, is given by

$$u_{x,t} = S \varepsilon_{x,t}, \quad (3)$$

with $S = A_x^{-1}$. The variance-covariance matrix of the reduced-form model Σ can be written as

$$E [u_{x,t} u'_{x,t}] = E [SS'] = \Sigma. \quad (4)$$

Now, let r be the interest rate included as one of the macroeconomic variables in this VAR block, and let $X_{q,t}$ be macroeconomic variables other than r . We denote reduced-form residuals $u_{x,t}$ of r as $u_{r,t}$ and of $X_{q,t}$ as $u_{q,t}$. We define $\varepsilon_{r,t}$ and $\varepsilon_{q,t}$ similarly to $u_{r,t}$

¹⁰We use 12 lags, $p = 12$, in our baseline estimation. The results are little changed when other numbers of lags, $p = 6$ or 18, are used for the estimation.

and $u_{q,t}$. Regarding the component of the matrix S , s is defined as the column in matrix S corresponding to $\varepsilon_{r,t}$. To obtain the impulse response to the monetary policy shocks, we do not have to identify all of the coefficients in S . Specifically, we need to estimate the following equation:

$$X_t = c_x + \sum_{j=1}^p B_{xx,j} X_{t-j} + s\varepsilon_{r,t}. \quad (5)$$

In order to obtain the coefficient of the matrix $B_{xx,j}$, we can simply use the least-squares estimation in Equation (2). Since we focus on the case of Japan where the policy target rate is constrained by the effective lower bound, it is important for a policy indicator to include shocks to unconventional monetary policy such as forward guidance. Thus, we choose one-year and two-year rates as candidates for the policy indicator. To isolate movements of the policy indicator made by exogenous monetary policy shocks, we use instrument variable Z_t . The conditions of the instrument variables Z_t are

$$E[Z_t \varepsilon_{r,t}] = \rho, \quad (6)$$

$$E[Z_t \varepsilon_{q,t}] = 0. \quad (7)$$

These equations state that instruments Z_t are only correlated with monetary policy shocks $\varepsilon_{r,t}$, but they are orthogonal to other macroeconomic variables shocks $\varepsilon_{q,t}$. If we are able to obtain these instrument variables, s can be estimated as follows.

- From the instrument variables Z_t , and the residual u_t obtained from the VAR regression, estimate the following equation:

$$u_{r,t} = \alpha_1 + \beta_1 Z_t + \xi_{1,t}, \quad (8)$$

where α_1 and β_1 are coefficients and $\xi_{1,t}$ is the residual. From the equation, we can calculate the fitted value $\hat{u}_{r,t}$.

- Using the fitted value $\hat{u}_{r,t}$, estimate the following equation using the ordinary least squares (OLS):

$$u_{q,t} = \alpha_2 + \beta_2 \hat{u}_{r,t} + \xi_{2,t}. \quad (9)$$

The estimated coefficient β_2 corresponds to the ratio s_q/s_r , where s_r and s_q are elements of s and they correspond to the response of

$u_{r,t}$ and $u_{q,t}$ to a unit increase in the policy shocks $\varepsilon_{r,t}$. Finally, s_r can be obtained from the reduced-form variance-covariance matrix.¹¹

Given these coefficients of s and $B_{xx,j}$, the impulse response to the monetary policy shocks can be estimated from the macrovariable block.

3.2 Segment-Variable Block

We denote the N_y -dimensional vector of segment variables at time t by Y_t . In this segment-variable block, the dynamics of macroeconomic variables do not depend on segment-level variables, while those of the segment variables depend on the macroeconomic variable block. Specifically, the reduced-form VAR can be described as follows:

$$Y_t = c_y + \sum_{j=0}^{q_1} B_{yx,j} X_{t-j} + \sum_{j=1}^{q_2} B_{yy,j} Y_{t-j} + u_{y,t}, \quad (10)$$

where $u_{y,t}$ is the reduced-form shocks of segment variables.¹² Here, we assume that the macroeconomic variables have contemporaneous effects on the segment variables. We can estimate the impulse response functions of the segment-variable block using the above relationship. Specifically, for the second term which corresponds to the behavior of the macroeconomic variables, we substitute an estimated impulse response function in the macroeconomic-variable block using the external instruments to obtain the impulse response functions.

The standard errors are calculated by the wild bootstrap following Mertens and Ravn (2013). The wild bootstrap generates valid confidence intervals under heteroskedasticity. We consider not only estimation errors related to macroeconomic and segment-level variables but also those related to instrument variables, by including all the steps of estimation for the bootstrap procedure.

¹¹Details of the estimation of s_r are described in Gertler and Karadi (2015).

¹²In the baseline estimation, we use $q_1 = q_2 = 4$ on a quarterly basis since some of the segment-level variables are only available in quarters. Even if we change the number of lags to shorter (1 and 2) or longer (6) values, we obtain similar results.

3.3 Assumptions of VAR Model

Our specification of the VAR model is based on several assumptions. First, we assume that the instrument variables using high-frequency data can extract the exogenous component of monetary policy shocks. In other words, the price changes within the narrow window between monetary policy announcements are only due to the monetary policy, and do not contain other economic or financial news. The data we use are described in the next section, and the validity of the assumption is checked in Nakamura, Sudo, and Sugisaki (2024).

Second, we assume that macroeconomic variables are isolated from the segment-variable block, and that the dynamics of macroeconomic variables do not depend on the segment variables. In order to check the validity of this block recursive restriction on the lag coefficients as we do here, we use the single-equation F-test, following Lee and Ni (2002). The results show that segment variables are not significant in the macroeconomic-variable equations, which justifies our assumption.

4. Data

Our data sample consists of three groups of variables: aggregate data, disaggregated labor-related variables by age and gender, and high-frequency data used for the identification of monetary policy shocks.

4.1 Macroeconomic Variables and Disaggregated Labor-Related Variables

The set of macroeconomic variables used for estimating the VAR in Equation (2) includes the one-year government bond yield, corporate bond yield, consumer price index (CPI), capacity utilization, unemployment rate, and number of employed. They are all used in the VAR in the log level, except for the one-year government bond yield, corporate bond yield, and the unemployment rate. Financial data such as the one-year rate and corporate bond yield come from Bloomberg. CPI, unemployment rate, and number of employed are taken from statistics released by the Ministry of Internal Affairs and

Communications, and the capacity utilization rate is taken from the Indices of Industrial Production released by the Ministry of Economy, Trade and Industry. The disaggregated labor-related variables, including the employment status by age and gender, are taken from the Labor Force Survey released by the Ministry of Internal Affairs and Communications. The sample period is from 1999:Q1 to 2018:Q4.¹³ The data series used for the estimation are shown in Figure 3.

4.2 High-Frequency Data

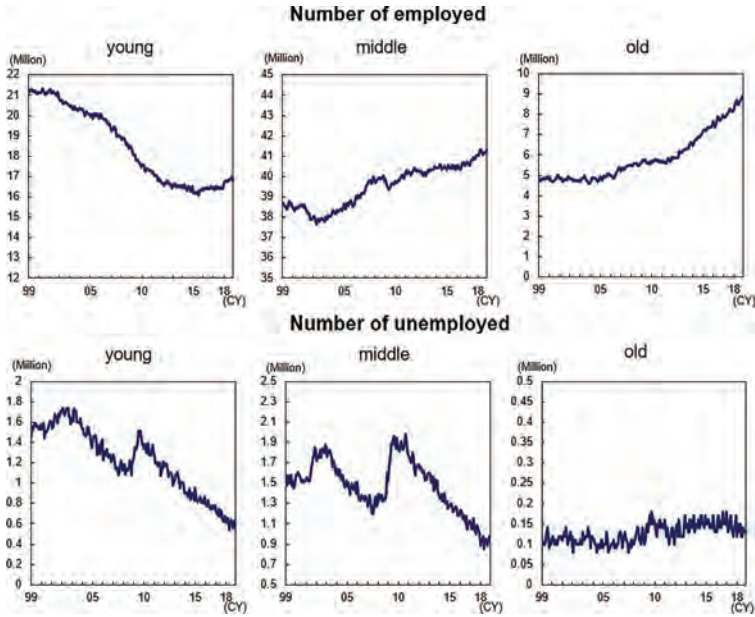
High-frequency data used as instrument variables are the key to our econometric framework.¹⁴ We use tick-by-tick price data of interest rate futures.¹⁵ The data we use are the three-month Euroyen interest rate futures (YE) with a variety of contract dates: 1 day to 3 months ahead (YE1), 3 to 6 months (YE2), 6 to 9 months (YE3), and 9 to 12 months (YE4). For the baseline estimation, following Nakamura, Sudo, and Sugisaki (2024), we use the primary factor extracted from changes in YE1 to YE4 (denoted YEF hereafter) over a 30-minute window around the announcement of the monetary policy meetings as the external instrument. These high-frequency data come from the Tokyo Financial Exchange Inc.¹⁶

¹³Disaggregated labor-related variables are only available from 2002:Q1 onward. When estimating the response of these variables (namely Figure 6), we use the sample period that runs from 2002:Q1 to 2018:Q4. Partly because of this short sample period, Figure 6 shows larger standard errors, and thus we use a longer sample period for other baseline results. Even if we use a shorter sample period when estimating the impulse response of the number of employed, unemployed, and the labor force, the baseline results remain consistent, albeit with larger standard errors.

¹⁴The identification of monetary policy shocks using high-frequency financial data was first developed in the U.S. by studies such as Gürkaynak, Sack, and Swanson (2005), Gertler and Karadi (2015), and Nakamura and Steinsson (2018).

¹⁵Other widely used methods for the identification of monetary policy shocks include Romer and Romer (2004). Studies such as Coibion et al. (2017) and Nakamura (2019) use such methods for the identification for the empirical analysis in the U.S. using the sample period before the Global Financial Crisis. However, the same technique is not applicable to the period that includes the period when effective lower bound prevails, and therefore is not applicable for Japan in our sample period.

¹⁶We follow Munakata, Oi, and Ueno (2019) for the data cleaning regarding the high-frequency data.

Figure 3. Data Used for the Estimation by Age Group

Note: The data come from the Labor Force Survey. Young, middle, and old represent age groups of under 35, between 35 and 64, and 65 and over, respectively.

To check the validity of the extracted monetary policy shocks, we estimate the response of the financial variables. Specifically, we use the following equation:

$$\Delta R_t = \alpha + \beta \Delta i_t + \varepsilon_t, \quad (11)$$

where ΔR_t and Δi_t correspond to the change in asset return and the change in interest rate on the day that a policy announcement is made. We estimate this equation using the two-stage least squares by instrumenting YEF for a daily change in a policy indicator, Δi_t . Our identifying assumption is that the instrument variables (YEF changes) are orthogonal to the error term, and the instrument affects ΔR_t only through the policy indicator Δi_t . We estimate the regressions over the available 2003:M4–2017:M10 samples.

Table 1 shows the results of the two-stage regression. Each row represents a particular policy indicator. The coefficient represents

Table 1. Effects of Monetary Policy Shocks on Financial Markets

Policy Indicator	2 Year	5 Year	10 Year	USD/JPY
1 Year	0.1211*** (4.235)	0.1457*** (5.377)	0.1457** (2.591)	-4.3118** (-2.069)
2 Year		0.1203*** (5.147)	0.1203** (2.512)	-3.5592*** (-2.762)
Note: Robust t-statistics in parentheses. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent level, respectively.				

the impact of a 10 basis point (bp) increase in a given policy indicator due to an exogenous monetary policy shock on a corresponding asset return. From the table, we can find that monetary policy easing shocks (in Japan) decrease 2- to 10-year government bond rates and depreciate the value of the Japanese yen, as expected. All of the results except for the 10-year bond rate and the exchange rate with 1-year government bonds as policy indicators are statistically significant at the 1 percent level, and all results are statistically significant at the 5 percent level.

As shown in Table 1, high-frequency data of monetary policy shocks indeed have a statistically significant impact on financial variables. This result indicates that YEF can be used as instruments when estimating the effects of monetary policy shocks. Then, the next issue is the choice of a policy indicator in the VAR specification: a financial variable whose shock corresponds to monetary policy innovation. As discussed in Gertler and Karadi (2015), traditional VARs use the same policy instruments and policy indicator, namely the overnight interest rate. This is because a structural shock to a policy indicator can be regarded as an exogenous monetary policy shock. On the other hand, we use the longer maturity rate of government bonds to capture shifts in unconventional monetary policy such as forward guidance. The longer-term government bond might incorporate various information other than monetary policy such as news on the economy reflecting the economic fundamentals. However, the high-frequency identification successfully eliminate the component of innovations in longer-term government bonds that is unrelated to monetary policy shocks.

**Table 2. Effects of Instruments
on the First-Stage Residuals**

Policy Indicator	1 Year	2 Year
Coefficient	1.61***	1.85***
R^2	0.07	0.08
Robust F-statistic	13.11	12.01
Note: *** denotes significance at the 1 percent level.		

When choosing the policy indicator, we need to avoid the weak instrument problem. Stock, Wright, and Yogo (2002) recommended that the F-statistic from the first-stage regression in the two-stage least squares should be above 10 in order to state that weak instrument problems are not present. Thus, we estimate the effects of high-frequency instruments on the first-stage residuals described in Equation (2). Table 2 shows the results. The left column shows the result when the one-year rate is the policy indicator, and the right column shows the case when the two-year rate is the policy indicator. The first row shows the estimates for the coefficient and the second row shows the R^2 . The third row shows the robust F-statistic for each regression. From the table, we can see that monetary policy shocks denoted as YEF explain more than 5 percent of the monthly innovation of the two-year rate, and the associated robust F-statistic is around 13, which is above the threshold value, 10. On the other hand, although the robust F-statistic for the two-year rate is more than 10, the value for robust F-statistics is less than it is in the case for the one-year rate. Thus, we use the one-year rate as a policy indicator in our baseline estimation.¹⁷

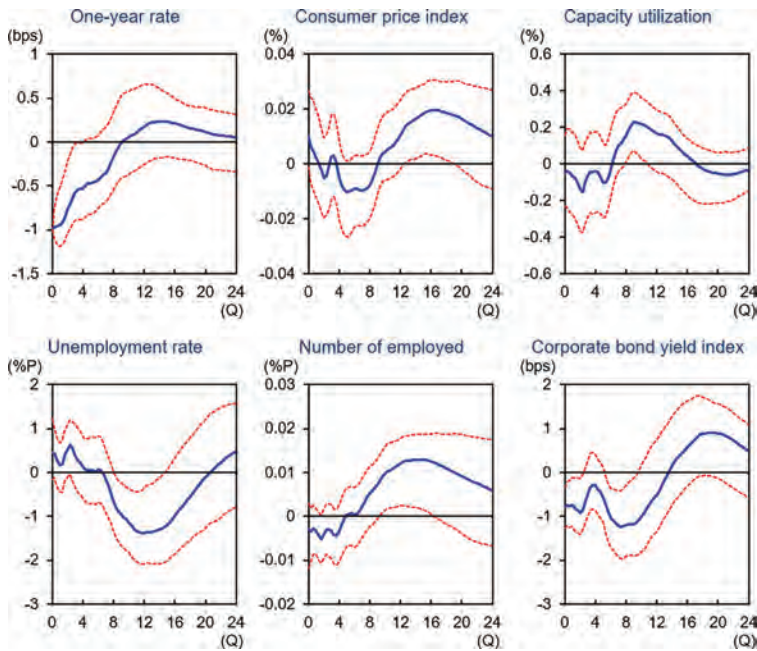
5. Empirical Results

5.1 Macroeconomic Variable Responses

Figure 4 shows the impulse response function of the macroeconomic variables to an expansionary monetary policy shock estimated using

¹⁷Even if we use the two-year rate as a policy indicator, we obtain almost the same results.

Figure 4. Macroeconomic Responses to a Monetary Policy Innovation



Note: This figure shows the responses of macroeconomic variables to an expansionary monetary policy shock. The shock is identified by high-frequency data of YEF and normalized so that the one-year rate decreases by 1 bp. Dashed lines show 90 percent confidence intervals.

the VAR in Equation (2). The size of the monetary policy shock is normalized so that the one-year rate declines by 1 basis point at the time of the impact, and the shock is identified by high-frequency data of YEF (factor of the Euroyen interest rate futures). Dashed lines show 90 percent confidence intervals. Consistent with the conventional view about the effects of an expansionary monetary policy shock, the CPI, capacity utilization, and number of employed all increase while the unemployment rate and corporate bond yield fall at a statistically significant level for all of the variables.¹⁸

¹⁸As for the variables used in the VAR analysis, we use the logarithmic scale in the CPI, capacity utilization, and number of employed for the estimation.

5.2 *Segment-Variable Responses*

5.2.1 *Number of Unemployed*

The upper charts of Figure 5 show the responses of the number of unemployed to an expansionary monetary policy shock by age. The number of unemployed does not respond in the first few quarters and starts to decline around two years after the shocks. The point estimates indicate that there is not much difference in terms of the timing and size of the decline across different age groups, while the declines of the unemployed for the old are muted, and the estimates are associated with larger confidence intervals. One potential reason why the response of the number of unemployed at and above 65 differs from that of the unemployment rate for other ages is the mandatory retirement age, which is now typically 60 in Japan.¹⁹ Due to this retirement, old individuals tend to become non-labor force rather than unemployed when they quit their current position.²⁰

5.2.2 *Number of Employed*

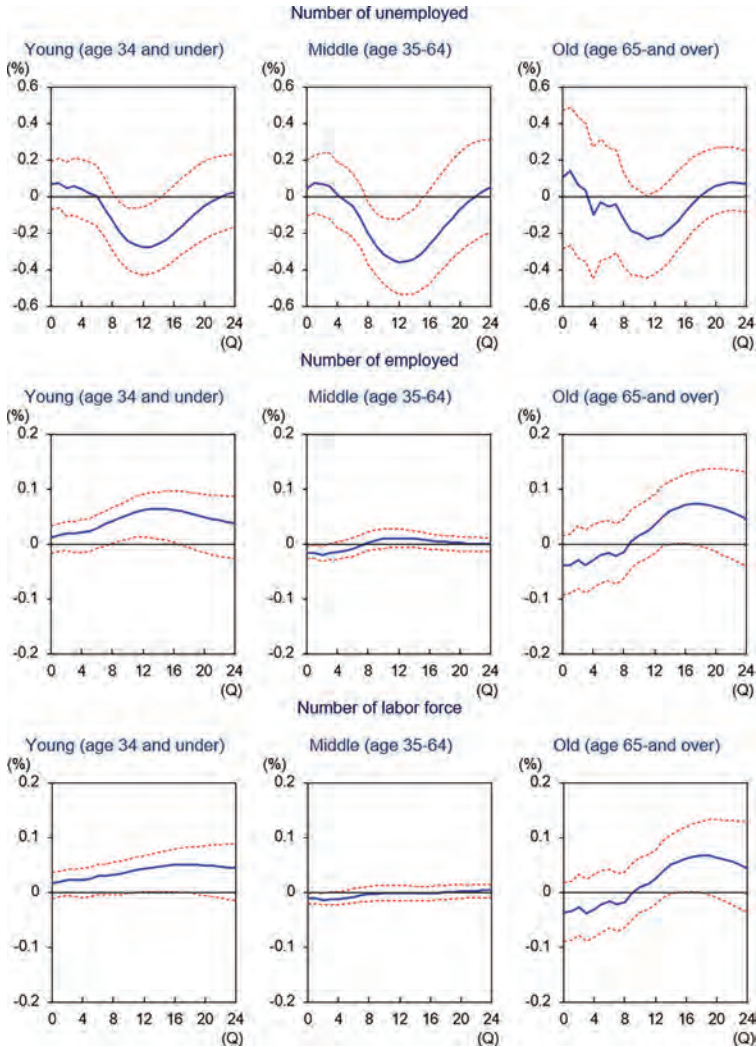
The middle charts in Figure 5 show the number of employed in response to an expansionary monetary policy shock.²¹ From these

¹⁹The Act on Stabilization of Employment of Elderly Persons, revised in 2000, required firms to make efforts to employ workers until the age of 65. The law was then revised in 2013, requiring firms to guarantee the employment of all personnel seeking to remain employed until the age of 65 by 2025.

²⁰We also check the impulse responses of the unemployment rate by age to an expansionary monetary policy shock, instead of the number of unemployed. We find that the results are unchanged qualitatively; however, we do not include them in this paper to conserve space. It is also possible to further examine the reason behind the changes in unemployment using the breakdown of unemployment reported in the Labor Force Survey. We decompose the unemployment into voluntary unemployment and involuntary unemployment. Note that involuntary unemployment refers to individuals who begin searching for a job due to factors attributable to their employer or business, whereas voluntary unemployment refers to those who leave their previous job for personal or family reasons. We find that the decline of the unemployment rate in response to an expansionary monetary policy shock is mainly due to the decrease in involuntary unemployment rather than voluntary unemployment.

²¹In this paper, the unemployed, the employed, and the total labor force population are divided into three mutually exclusive groups, and separate estimates are conducted for each group. On the other hand, the working paper version of this study employs a rolling sample approach to divide the

Figure 5. Employment Responses by Age



Note: This figure shows the responses of the number of unemployed, the number of employed, and the labor force to an expansionary monetary policy shock depending on the age given at the top of each graph. The shock is identified by high-frequency data of YEF and normalized so that the one-year rate decreases by 1 bp. Dashed lines show 90 percent confidence intervals.

data into eight overlapping groups. The results indicate no substantial differences in age-specific responses to monetary policy shocks between the three-group and eight-group classifications. See Figures 6–8 in the following link: <https://www.imes.boj.or.jp/research/abstracts/english/21-E-06.html>.

charts, we find that the increase in response to an expansionary monetary policy shock is large in younger age groups and older age groups. More specifically, the number of employed young workers, who are under age 46, and the number of those age 65 and over increase.²² On the other hand, the response of workers between the ages of 35 and 64 is close to zero over the six years after the shock.²³

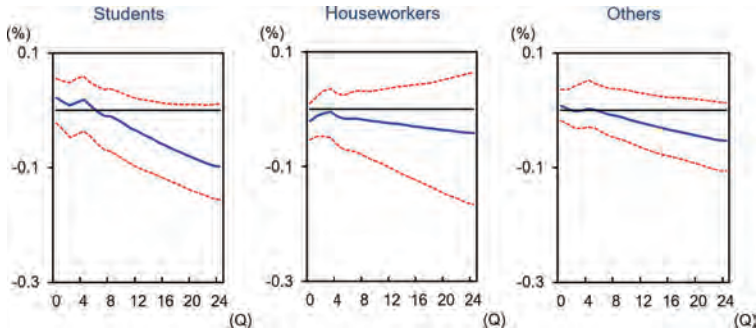
5.2.3 *Labor Force*

Generally, an increase in the number of employed occurs due to two reasons: a decrease in the number of unemployed, or an increase in the labor force. In our case, the former is unlikely to be the driving force, since the homogeneous impulse responses of the number of unemployed cannot account for the heterogeneous impulse responses of the number of employed. By contrast, as shown in the lower charts of Figure 5, we find evidence that the latter is the case: the age-specific shape of the response of labor force and that of the number of employed to a monetary policy shock are somewhat close. The size of the expansionary monetary policy shock is the same as in the previous figures. The figure reveals that young and old workers show larger responses. On the other hand, the response of workers between the ages of 35 and 64 does not increase after an expansionary monetary policy shock in the point estimate.

An increase in the number of labor force means that the number of non-labor force decreases. To explore the reason for this decrease, we decompose the non-labor force into students, homemakers, and others, following the definition in the Labor Force Survey. Figure 6

²²By estimating which type of worker contributes to the increase in the number of employed following an expansionary monetary policy shock, we find that, for the younger worker, the main driving force of the increase is regular employment. On the other hand, non-regular employment including part-time job is the main contributor to the increase in the number of employed for the elderly.


²³A large positive response of the youth following expansionary monetary policy shocks may contribute to alleviating the so-called scarring effect presented in, for example, Arulampalam, Booth, and Taylor (2000) and Arulampalam (2001) for the U.K. They argue that an individual's previous unemployment experience increases the probability of current unemployment and contributes to wage inequality or poverty. With respect to this scarring effect, our results show that expansionary monetary policy shocks increase the employment of the youth, which may reduce the risk of future unemployment, as younger workers accumulate human capital in their early stage of the work experiences.

Figure 6. Non-labor Force Responses

Note: This figure shows the responses of the non-labor force to an expansionary monetary policy shock. The shock is identified by high-frequency data of YEF and normalized so that the two-year rate decreases by 1 bp. Dashed lines show 90 percent confidence intervals.

shows the impulse response of each component to an expansionary monetary policy shock. Although all the responses are not statistically significant at the 90 percent confidence interval, the point estimates indicate that the size of the decrease is the largest in students, and larger in others compared to homemakers. Figure 7 shows the breakdown of the number of non-labor force by age. It shows that the number of population that is not in labor force is high among individuals ages 15 to 24, as well as those over 65, particularly. Moreover, students constitute the dominant component of the non-labor force between the ages of 15 and 24. For ages between 30 to 54, most of the non-labor force consists of homemakers, and for the old, especially those older than 65, the percentage of others is large.

Figures 6 and 7 suggest the reasons behind the heterogeneous response of the labor force to monetary policy shocks. First, it is notable that among the youth, students constitute a significant segment of the non-labor force and the response of the youth is considered to be related to the observations made in existing studies that the labor force participation of the younger population tends to be procyclical as they transition between students and workers; economic upturns encourage more students to enter the labor market, while downturns deter such engagement (see Kondo 2007). These dynamics are partially attributed to the

Figure 7. Breakdown of Non-labor Force

Note: This figure shows the breakdown of the non-labor force by age. The data come from the Labor Force Survey.

simultaneous recruitment of new graduates and the lifetime employment system, where companies hire fresh graduates and retain them until retirement. Consequently, for younger generations, the initial job—secured immediately after graduation—is of considerable importance for their lifetime employment prospects, as pointed out by Oshio and Inagaki (2015). Thus, during recessions when fewer regular employment opportunities are anticipated, more students in Japan opt to continue their education rather than entering the labor force in pursuit of desirable positions.

Second, for older individuals, there is an increase in the proportion of those categorized as “others,” including retirees making up a substantial part of the non-labor force. Their response is considered to be related to the observation that their labor force participation tends to be procyclical as they transition between workers and retirees. They may opt to reenter the workforce once the economy shows signs of improvement following an expansionary monetary policy shock. Given Japan’s high life expectancy, aged individuals often believe that the benefits from the Japanese pension system may not suffice to cover their living expenses. Additionally, given the particularly extended life expectancy in Japan, all else being equal, the level of consumption for retired households could likely be diminished. These factors may contribute to a sense of financial

insecurity among the Japanese elderly.²⁴ In summary, Japan's high life expectancy coupled with concerns that the benefits from the Japanese pension system may not adequately cover living expenses may have made Japanese elderly more responsive to labor market conditions compared with the case otherwise. This channel could work, reinforcing an outward shift of the labor supply curve when an expansionary monetary policy shock improves labor market conditions. Conversely, for those between the ages of 30 and 50, the majority of the non-labor force consists of homemakers, many of whom may be unable to seek employment even in an improved labor market condition due to childcare responsibilities.

5.2.4 *Gender Differences*

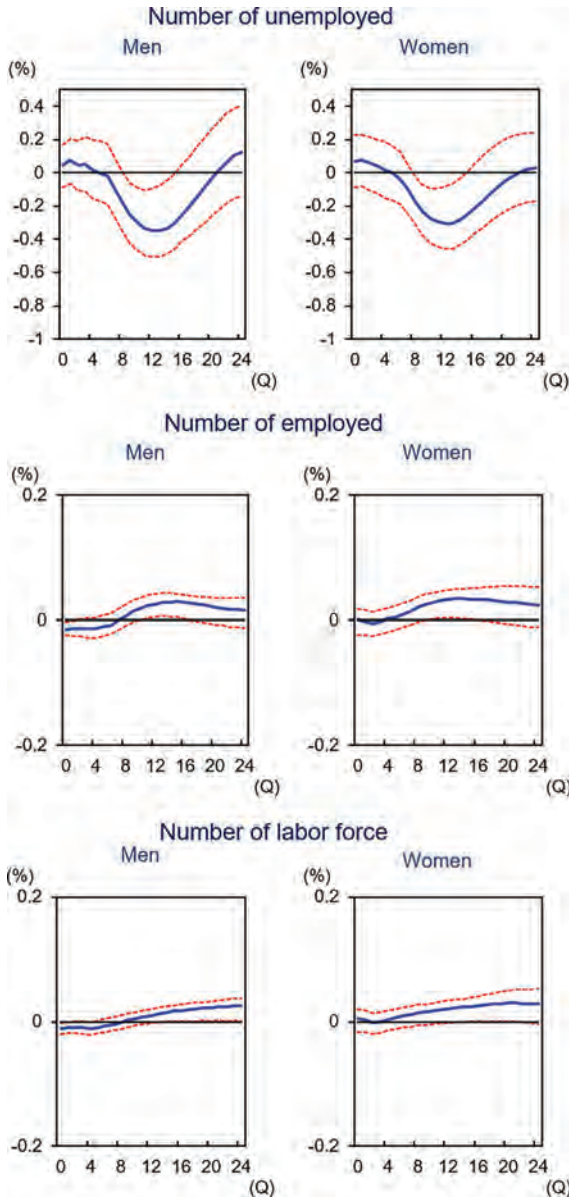
We also estimate the impulse response for each gender to check if the effects of monetary policy differ across gender. Figure 8 shows the responses for men and women to an expansionary monetary policy shock. Again, the shock is normalized so that the one-year rate decreases by 1 bp. Dashed lines show 90 percent confidence intervals.

From these results, we can conclude that the responses of employment, unemployment, and labor force are similar between men and women. Specifically, the chart clearly shows that an expansionary monetary policy shock reduces the number of unemployed, whereas it increases the number of employed and labor force statistically significantly. However, as we can see from the impulse response, the difference across the genders is not significant,²⁵ *inter alia*, compared

²⁴This sentiment is corroborated by a survey conducted by the Cabinet Office of Japan. It reveals that a significant portion of the elderly do not feel financially secure, in contrast to their counterparts in the U.S. and European countries: over half of the elderly in Japan perceive their wealth as insufficient, whereas this figure is around or below 20 percent in the U.S. and Europe. Consequently, about 40 percent of the elderly express a desire to work, primarily due to inadequate financial resources—a proportion substantially higher than that in peer countries.

²⁵Note that while the Bank of Japan (2018) indicates that the labor supply of women is more elastic than that of men, this paper finds that the response of the labor force to our identified monetary policy shocks is quite similar between the two genders. The Bank studies differences in the elasticity of labor supply reflecting various changes in the economic environment, including the Japanese government's policies that encourage the participation and advancement of

Figure 8. Responses by Gender



Note: This figure shows the responses of men and women to an expansionary monetary policy shock. The shock is identified by high-frequency data of YEF and normalized so that the one-year rate decreases by 1 bp. Dashed lines show 90 percent confidence intervals.

to the difference of impulse responses across age groups.²⁶ Also, to compare with the examples of other countries, results in the U.S. are explained in the appendix.

6. Implications for the Wage Phillips Curve

In the previous section, we show that the responses of the number of employed and the labor force to expansionary monetary policy shocks are larger among the young and old generations. This section studies an implication of these heterogeneous responses for nominal wage dynamics, using the concept of the wage Phillips curve.²⁷ Theoretically speaking, as Galí (2011) points out, the wage Phillips curve is flatter when the labor supply is elastic. Therefore, if the young and the old adjust their labor force inflow in response to a shock (as empirically shown in the previous section) to a greater degree, the wage Phillips curve for the young and the old, *ceteris paribus*, should be flatter than that for the middle-aged.²⁸

In order to assess this argument quantitatively, we estimate the wage Phillips curve for each age group of workers. For this estimation, we need the unemployment rate and wage inflation rate of each age group. The data for unemployment rate are taken from the Labor Force Survey published by the Ministry of Internal Affairs and Communications. We use the deviation from the average unemployment rate in each group to estimate the slope of the wage Phillips

women in the labor force, whereas our paper investigates changes in the number of employed and labor force in response to monetary policy shocks identified using high-frequency market data.

²⁶Even if we show the impulse response across the age decomposed by gender, the difference of the impulse response in gender is smaller compared to those of ages. We also check the response of unemployment rate instead of the number of unemployed, but the results are almost the same: the size of the reduction of the unemployment rate is almost the same between genders. These results are shown in the working paper version in the following link: <https://www.imes.boj.or.jp/research/abstracts/english/21-E-06.html>.

²⁷The derivation of the wage Phillips curve is explained in, for example, Erceg, Henderson, and Levin (2000) and Galí (2011). The wage Phillips curve in Japan is estimated, for instance, by Muto and Shintani (2020) and Iwasaki, Muto, and Shintani (2021).

²⁸Muto and Shintani (2020) point out that the slope of the wage Phillips curve is determined also by wage stickiness. Nonetheless, since wages for workers in Japan, regardless of age, are likely to be simultaneously revised in April by Shunto, it is reasonable to assume that wage stickiness is more or less the same across the ages.

curve.²⁹ The wage data come from the Basic Survey on Wage Structure published by the Ministry of Health, Labour and Welfare. Wage inflation rate is calculated from the baseline wage released in the Basic Survey on Wage Structure, scheduled salary in June excluding overwork salary.³⁰ The sample period is from 1999 to 2019.

Figure 9 shows the wage Phillips curve for each indicated age. Each circle corresponds to each data point, and the line shows the fitted curve obtained from the OLS in each age group. Also, we use the data of males and females separately to increase the number of data.³¹ The horizontal axis corresponds to the deviation of the unemployment rate from the average. The vertical axis corresponds to the wage inflation rate. From the figure, we can see that all of the slopes are negative; a decrease in unemployment leads to a higher wage inflation. Furthermore, most importantly, the slope of the wage Phillips curve of middle-aged workers is the steepest.

In order to check the statistical significance, we use two regression equations. The first one is the simple regression,³²

$$\pi_{w,t} = \alpha_0 + \alpha_1 unemp_t + \xi_t, \quad (12)$$

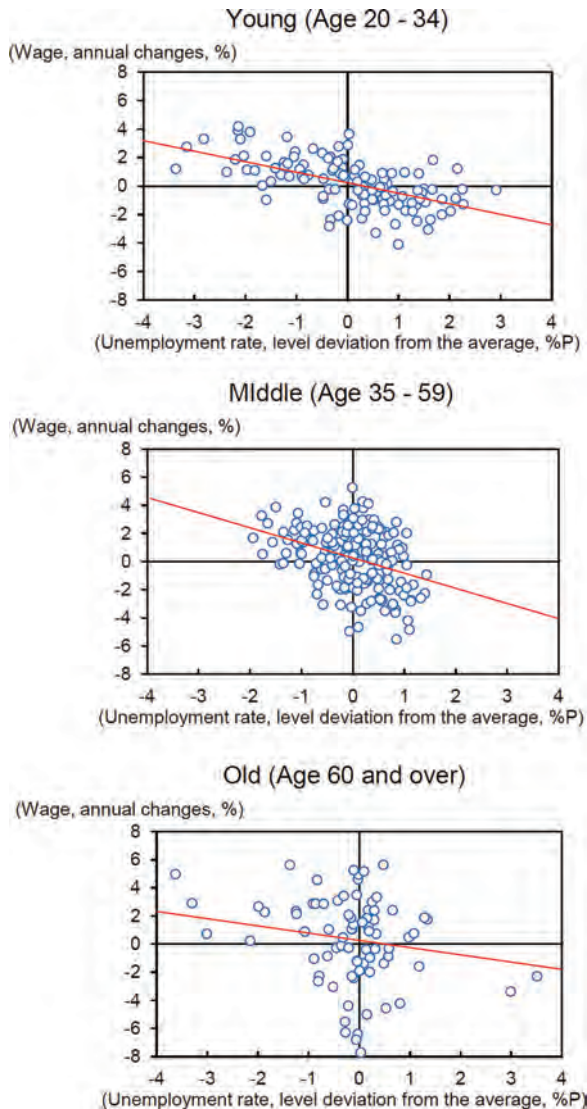
where $\pi_{w,t}$ is the wage inflation rate and $unemp_t$ is the unemployment rate. The α terms represent the regression coefficient, and ξ_t

²⁹We also detrend using a linear trend of time series of unemployment rate in each age. However, the method of calculating the deviation does not significantly change the results.

³⁰Another definition of wage inflation rate used in estimating the wage Phillips curve is monthly total wage divided by monthly total working hours. We also estimate the results based on this definition, and obtain qualitatively similar results. Moreover, the hourly wage of short-time workers also gives similar results: the slope of the middle-aged worker is the steepest.

³¹Since we use the Basic Survey on Wage Structure, only the wage change on a yearly basis is available. Thus, we use male and female data separately to increase the number of data. To ensure robustness, we conduct the same exercises using the consolidated data and, again, obtain qualitatively similar results to the baseline ones.

³²By assuming a simple autoregressive model for the determination of the unemployment rate, we can obtain the closed-form representation of the wage Phillips curve similar to Equation (12) rather than the standard forward-looking form of the New Keynesian wage Phillips curve, as described in Galí (2011). One may think that this simple specification cannot avoid the endogeneity problem. Indeed, several preceding studies such as Galí (2011) and Muto and Shintani (2020) control for inflation in estimating the wage Phillips curve. We also control this inflation rate as a robustness check, and obtain similar results; the wage Phillips curve for the young and the old is flatter.

Figure 9. The Wage Phillips Curve of Each Age Group

Note: This figure shows the wage Phillips curve of each age given at the top of each graph. The horizontal axis corresponds to the unemployment rate deviations from the average unemployment rate. The vertical axis corresponds to the wage inflation rate. Each circle corresponds to each data point, and the red lines show the fitted line from the OLS. The unemployment data come from the Ministry of Internal Affairs and Communications, and wage data come from the Ministry of Health, Labour and Welfare. The sample period is from 1999 to 2019.

Table 3. Wage Phillips Curve Depending on Age

Age	Unemployment Rate	Intercept
20–34 (Robust-t) (90% CI)	–0.74*** (–7.96) [–0.89 –0.59]	0.24* (1.95)
35–59 (Robust-t) (90% CI)	–1.08*** (–6.09) [–1.37 –0.78]	0.26* (1.96)
60– (Robust-t) (90% CI)	–0.51*** (–2.59) [–0.84 –0.19]	0.27 (0.87)

Note: ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent level, respectively. Robust-t denotes robust t-statistics. 90% CI denotes 90 percent confidence interval.

is the regression error term. This equation corresponds to the most basic wage Phillips-curve estimation, and we estimate one for each age group. Table 3 shows the regression results of a wage Phillips curve for each age group. The table also shows the robust t-statistics and 90 percent confidence interval. These regression results are used to describe the fitted line in Figure 9. The results indicate that the slope of the coefficient is statistically significant at the 1 percent level for all age groups, and the slope of middle age is the steepest. Moreover, the point estimate of middle age is out of the 90 percent confidence interval of young and old workers.³³

The second estimation is the pooled regression with age dummy variable written as

$$\pi_{w,t} = \alpha_0 + \alpha_1 unemp_t + \alpha_{2,x} dummy_{x,t} unemp_t + \xi_t. \quad (13)$$

³³In this estimation of the wage Phillips curve, we use a slightly different age group compared to the VAR analysis. While we maintain the same threshold between young and middle-aged workers, the threshold between middle-aged and older workers is 65 for VAR analysis, whereas it is 60 for wage Phillips-curve estimation. This adjustment is due to the limited availability of wage data for older workers, which results in an insufficient sample size for estimating the wage Phillips curve if we use 65 years as the threshold. Consequently, it becomes infeasible to employ the exact same threshold as in the VAR analysis.

Table 4. Estimation of the Wage Phillips Curve with Dummy Variables

Age	u	u*dmy(20-34)	u*dmy(35-59)	u*dmy(60-)	Intercept
Case 1	-0.74*** (-7.19)				0.25*** (2.58)
Case 2	-0.75*** (-4.10)	0.01 (0.07)			0.25*** (2.59)
Case 3	-0.66*** (-6.34)		-0.42** (-2.14)		0.26*** (2.64)
Case 4	-0.84*** (-9.23)			0.32 (1.42)	0.26*** (2.58)
Case 5	-1.08*** (-6.74)	0.34* (1.82)		0.56** (2.17)	0.26*** (2.60)

Note: ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent level, respectively. Robust t-statistics in parentheses.

Here, $dummy_{x,t}$ corresponds to the dummy variable, which takes a value of 1 at each age group described in Table 4. Table 4 shows the results. It reveals that all the coefficients for unemployment are statistically significant at the 1 percent level. Moreover, the point estimate of unemployment rate times dummy (35–59) is negative, while that with dummy (20–34) and dummy (60–) is positive. More importantly, in case 3, α_2 is statistically significant at the 5 percent level when dummy (35–59) is used; together with that in case 5, α_2 is significantly different from zero when dummy (20–34) and dummy (60–) are used simultaneously. These results confirm that the slope of the wage Phillips curve is the steepest for middle-aged workers, which is consistent with the VAR analysis obtained in the previous section.³⁴

7. Conclusion

In this paper, we address the question of which age group of workers' employment, unemployment, and labor force is the most responsive to monetary policy shocks, and try to find the macroeconomic

³⁴In addition, the appendix shows the results of estimating the wage Phillips curve in the U.S. We find that the slopes of the wage Phillips curve are similar across ages. This is consistent with the impulse response of the U.S. in the sense that the inflow of labor force does not differ across ages.

implications. Using the block recursive structure of the VAR model as well as monetary policy shocks identified by high-frequency data in Japan, we find that the responses of the number of employed to the shocks are large among young and old workers compared to the middle-aged workers, and the similar observations hold for the number in the labor force. On the other hand, the unemployment rate shows homogeneous impulse responses across all the age groups. These results suggest that there exists an added degree of underutilization of young and old workers before expansionary monetary shocks occur, which, other things being equal in turn, leads to the weak responses of wage inflation following shocks. Indeed, we find that the wage Phillips curve of middle-aged workers is steeper than those of young and old workers.

Appendix. Estimation Using U.S. Data

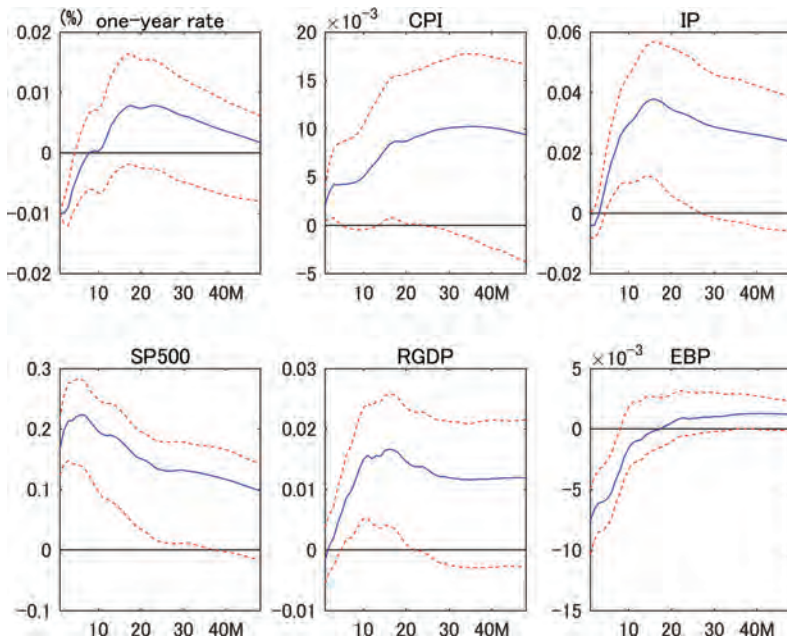
In this appendix, we present empirical results for the U.S., utilizing the same econometric framework as outlined in Section 3.

First, regarding the high-frequency data, which is used as the external instruments, we use the shocks in the three-month-ahead futures rate from Gertler and Karadi (2015). The one-year U.S. government bond rate is used as the policy indicator for the baseline analysis, but using the two-year bond rate gives very similar results. For the macroeconomic variables, we use the CPI, industrial production (IP), Standard and Poor's 500 index (SP500), real gross domestic product (RGDP), and excess bond premium (EBP), which is defined in Gilchrist and Zakrajšek (2012).³⁵ Regarding the segment variable of employment, we use the Current Population Survey (CPS) for the specific age groups of individuals over 15 years of age and older. The sample period is from 1979:M7 to 2016:M12.³⁶ The

³⁵We use macroeconomic-variables data from Jarociński and Karadi (2020). In this data set, GDP is interpolated monthly following Stock and Watson (2010). Basically, it uses a Kalman filter to distribute the quarterly GDP across months using a data set of monthly variables that are closely related to economic activity. The robust F-statistic when estimating the effects of high-frequency instruments on the first-stage residuals is 12.7, which is above 10. This indicates that the weak instrument problem is not present.

³⁶We use the 1979:M7 to 2016:M12 samples for the estimation of VAR coefficients. For the instruments variable, which is used to identify the monetary

Figure A.1. Macroeconomic Variable Responses in the U.S.

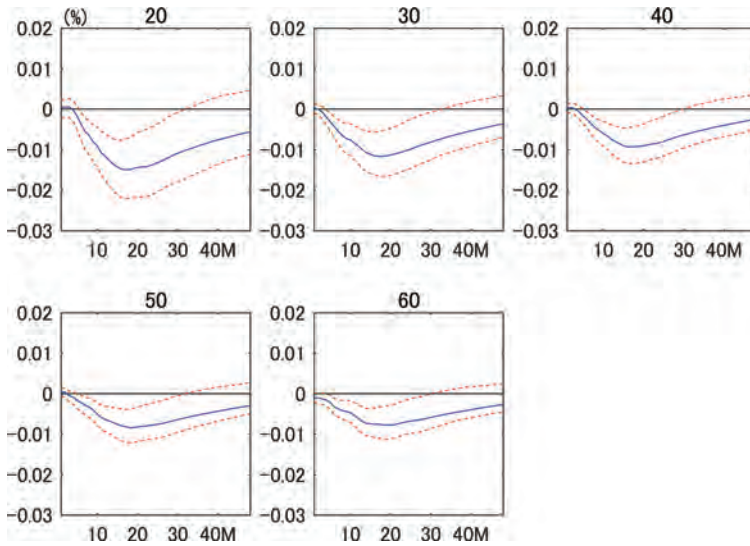


Note: This figure shows the responses of macroeconomic variables to an expansionary monetary policy shock in the U.S. The shock is normalized so that the one-year rate decreases by 1 bp. Dashed lines show 90 percent confidence intervals.

starting point coincides with the beginning of Paul Volker's tenure as Chairman of the Federal Reserve.

Figure A.1 shows the macroeconomic-variable responses to an expansionary monetary policy shock in the U.S. The shock is normalized so that the one-year bond rate decreases by 1 bp. We can see that in response to an expansionary monetary policy shock, the CPI, IP, and RGDP improve and SP500 rises, while the EBP, which is

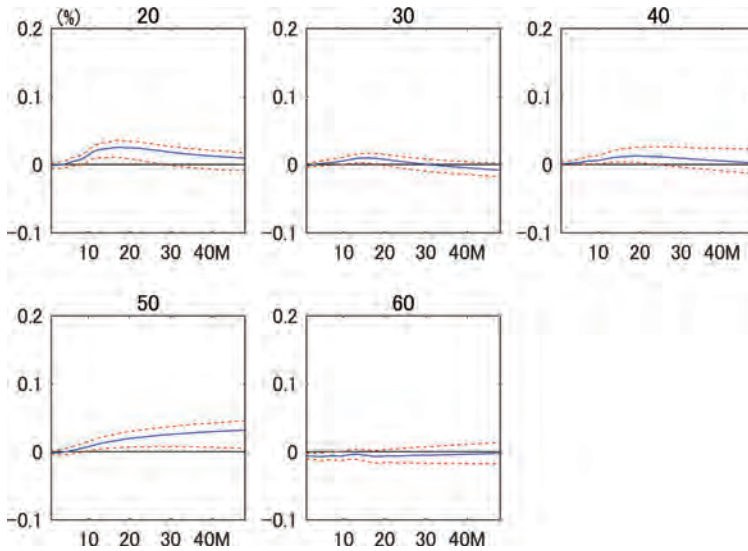
policy shocks, we use the available 1991:M1–2012:M6 sample for the baseline two-step regression. We estimate the results excluding 2008:M7–2009:M6, which corresponds to the period of financial turbulence, but the results are very similar to the ones including the financial turbulence period.

Figure A.2. Unemployment Rate Responses in the U.S.

Note: This figure shows the responses of the unemployment rate to an expansionary monetary policy shock at the age given at the top of each graph. The shock is normalized so that the one-year rate decreases by 1 bp. Dashed lines show 90 percent confidence intervals.

an indicator of credit spreads, falls. The response is consistent with the high-frequency identification VAR analysis using U.S. macroeconomic data such as Gertler and Karadi (2015) and Jarociński and Karadi (2020).

Figure A.2 shows the responses of the unemployment rate to an expansionary monetary policy shock by age (given at the top of each graph). Each graph at age x is estimated using the sample between age $x - 5$ and $x + 5$. Dashed lines show 90 percent confidence intervals. Importantly, the responses of the unemployment rate are homogeneous across ages. On the other hand, Figure A.3 shows the results for the number of employed. The figure shows that the response of the number of employed in the U.S. is large at age 20 and age 50. Compared to these ages, each of the responses at age 30 and age 40 is weaker. Figure A.4 shows the results of labor force response. We can see that each of the responses at age 20 and age 50 is also larger compared to that


Figure A.3. Number of Employed Responses in the U.S.

Note: This figure shows the responses of the number of employed to an expansionary monetary policy shock at the age given at the top of each graph. The shock is normalized so that the one-year rate decreases by 1 bp. Dashed lines show 90 percent confidence intervals.

of the middle-aged, which is similar to the number of employed case.³⁷


Compared to the results of the estimation in Japan, we find a similar pattern in the results for the U.S. in two respects. First, the responses of the unemployment rate to a monetary policy shock are homogeneous across the age groups. Second, the number of employed and labor force shows heterogeneous responses across ages including

³⁷Compared to the empirical results in Japan, in which the response of the number of employed and that of labor force show similar results, the response of labor force in the U.S. is smaller overall than that of employed. This might be because the size of the response of employed is smaller than that in Japan, which means that a decrease in the number of unemployed contributes to an increase in the number of employed. Supporting this view, the response of the labor force is smaller than that of employed between 10 to 20 months after the shock, which corresponds to the period when the response of the unemployment rate is the smallest.

Figure A.4. Labor Force Responses in the U.S.

Note: This figure shows the responses of the labor force to an expansionary monetary policy shock at the age given at the top of each graph. The shock is normalized so that the one-year rate decreases by 1 bp. Dashed lines show 90 percent confidence intervals.

the pronounced increase in the people of their 20s and 50s for the number of employed and in the people of their 50s for the labor force. The difference between Japan and the U.S. is the response of the number of employed for middle-aged workers: a larger response can be observed in the U.S. results. This might arise from the difference in labor market structure between Japan and the U.S. As Owan (2004) points out, labor mobility is lower in Japan than in the U.S. Figure A.5 supports this argument. It shows the years of tenure with the current employer in Japan and the U.S. by age group. Older workers tend to have worked longer in their current job, but the key point is that workers in Japan show longer years of tenure with their current employer. For example, between ages 35 and 44, the job tenure in Japan is about 10 years, while it is 5 years in the U.S. Meanwhile, wages tend to be lower for young and old workers in the U.S. as well, as shown in Figure A.6. For all age groups, job tenure in


Figure A.5. Years of Tenure with Current Employer

Note: This figure shows the years of tenure with the current employer in Japan and the U.S. for each age group in 2018. Japanese data come from the Ministry of Health, Labour and Welfare. The U.S. data come from the Bureau of Labor Statistics.

Japan is double that in the U.S.³⁸ Figure A.7 shows the number of job transfers and the job transfer rate by age group in Japan. More specifically, it shows the number of employed persons who changed jobs in the past one year (left axis, ten thousand), and the rate of employed persons who changed jobs in the past one year (right axis, percent). We can see that the youngest age group (between 15 and 24) shows a job transfer rate as high as 12 percent, while for older ages it is low, and above age 35 it is less than 5 percent. In sum, the job market is more liquid in the U.S. than it is in Japan, and the small response of the number of employed for the middle-aged workers to an expansionary monetary policy shock might be attributed to the low labor mobility in Japan, since middle-aged workers have


³⁸We need to note that the definition of tenure is slightly different between Japan and the U.S. Specifically, in Japan, employee tenure is a measure of how long workers have been with their current employer, while in the U.S. it is a measure of how long wage and salary workers have been with their current employer at the time of the survey. Thus, employee tenure may be shorter in the U.S. However, promotion within the company is not so common in the U.S. compared to Japan, and thus this difference of definition might not significantly change the results.

Figure A.6. Weekly Earnings by Age in the U.S.



Note: This figure shows the median weekly earnings in 2021 in the U.S. The data come from the U.S. Bureau of Labor Statistics.

Figure A.7. Job Transfer Rate in Japan



Note: This figure shows the employed persons who changed jobs in the past one year (left axis, ten thousand), and the rate of employed persons who changed jobs in the past one year (right axis, percent). The data come from the Labor Force Survey.

relatively fewer job vacancy opportunities even if business conditions improve.

We estimate the wage Phillips curve for each age group of workers in the U.S. The unemployment data come from the Current Population Survey (CPS). The wage data come from the median usual

Table A.1. Wage Phillips Curve in the U.S.

Age	Unemployment Rate	Intercept
20	-0.11*	0.75***
(Robust-t)	(-1.71)	(3.81)
30	-0.16***	0.69***
(Robust-t)	(-4.54)	(10.59)
40	-0.11***	0.68***
(Robust-t)	(-3.82)	(14.36)
50	-0.10**	0.62***
(Robust-t)	(-2.62)	(10.84)

Note: ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent level, respectively. Robust-t denotes robust t-statistics.

weekly earnings full-time wage.³⁹ The sample period is from 1996 to 2019.⁴⁰

In order to check the statistical significance, we use two regression equations, the same as the regression in the main text. The first one is the simple regression,

$$\pi_{w,t} = \alpha_0 + \alpha_1 unemp_t + \xi_t, \quad (\text{A.1})$$

where $\pi_{w,t}$ is the wage inflation rate and $unemp_t$ is the unemployment rate. The α terms represent the regression coefficient, and ξ_t is the regression error term. Table A.1 shows the regression results for the wage Phillips curve of each age group. The results indicate that the slope is similar across the ages.

The second estimation is the pooled regression with the age dummy variables,

$$\pi_{w,t} = \alpha_0 + \alpha_1 unemp_t + \alpha_{2,x} dummy_{x,t} unemp_t + \xi_t, \quad (\text{A.2})$$

³⁹Galí (2011) uses the average hourly earnings from the Establishment Survey, which is from the company data. On the other hand, we need the age information and use the data from the household side.

⁴⁰To estimate the slope of the wage Phillips curve at age x , we use the age from $x - 5$ to $x + 4$.

Table A.2. Estimation of the Wage Phillips Curve with Dummy Variables in the U.S.

Age	u	u*dmy20	u*dmy30	u*dmy40	u*dmy50	Intercept
Case 1	-0.12** (-2.42)	0.01 (0.10)				0.71*** (8.61)
Case 2	-0.11** (-2.08)		-0.05 (-0.48)			0.71*** (8.61)
Case 3	-0.11*** (-3.29)			0.00 (0.00)		0.71*** (8.61)
Case 4	-0.11*** (-2.26)				0.01 (0.12)	0.71*** (8.61)

Note: ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent level, respectively. Robust t-statistics in parentheses.

where $dummy_{x,t}$ takes a value of 1 at age x .⁴¹ For this regression, we use the sample between ages 16 and 54. Table A.2 shows the results. These show that none of the coefficients of the dummy variables is statistically significant. The results are consistent with the impulse response described in the main text. Namely, the impulse responses of the labor force in the U.S. are heterogeneous, but the differences between ages are not so distinctive compared to the results in Japan. Thus, the slope of the wage Phillips curve does not change with age in the U.S. regression.

References

- Abel, A. B. 2003. "The Effects of a Baby Boom on Stock Prices and Capital Accumulation in the Presence of Social Security." *Econometrica* 71 (2): 551–78.
- Arulampalam, W. 2001. "Is Unemployment Really Scarring? Effects of Unemployment Experiences on Wages." *Economic Journal* 111 (475): 585–606.
- Arulampalam, W., A. L. Booth, and M. P. Taylor. 2000. "Unemployment Persistence." *Oxford Economic Papers* 52 (1): 24–50.

⁴¹Strictly speaking, the dummy variable takes a value of 1 when the data are from age $x - 5$ to $x + 4$.

- Bahaj, S., A. Foulis, G. Pinter, and P. Surico. 2019. "Employment and the Collateral Channel of Monetary Policy." Working Paper No. 827, Bank of England.
- Bank of Japan. 2018. "The Recent Increase in Labor Supply and Wage Developments." *Outlook for Economic Activity and Prices* (July): Box 1, 42–44.
- . 2023. "Current Situation and Outlook for Labor Market Conditions." *Outlook for Economic Activity and Prices* (January): Box 2, 51–53.
- Bernanke, B. S. 2015. "Monetary Policy and Inequality." *Ben Bernanke's Blog*, June 1. <https://www.brookings.edu/blog/ben-bernanke/2015/06/01/monetary-policy-and-inequality/>.
- Blanchflower, D. G., and A. T. Levin. 2015. "Labor Market Slack and Monetary Policy." NBER Working Paper No. 21094.
- Cairo, I., S. Fujita, and C. Morales-Jimenez. 2022. "The Cyclicity of Labor Force Participation Flows: The Role of Labor Supply Elasticities and Wage Rigidity." *Review of Economic Dynamics* 43 (January): 197–216.
- Campolmi, A., and S. Gnocchi. 2016. "Labor Market Participation, Unemployment and Monetary Policy." *Journal of Monetary Economics* 79 (May): 17–29.
- Chodorow-Reich, G., and L. Karabarbounis. 2016. "The Cyclicity of the Opportunity Cost of Employment." *Journal of Political Economy* 124 (6): 1563–1618.
- Christiano, L. J., M. Trabandt, and K. Walentin. 2021. "Involuntary Unemployment and the Business Cycle." *Review of Economic Dynamics* 39 (January): 26–54.
- Cloyne, J., C. Ferreira, and P. Surico. 2020. "Monetary Policy when Households Have Debt: New Evidence on the Transmission Mechanism." *Review of Economic Studies* 87 (1): 102–9.
- Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia. 2017. "Innocent Bystanders? Monetary Policy and Inequality." *Journal of Monetary Economics* 88 (June): 70–89.
- 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.
- Erceg, C. J., D. W. Henderson, and A. T. Levin. 2000. "Optimal Monetary Policy with Staggered Wage and Price Contracts." *Journal of Monetary Economics* 46 (2): 281–313.

- Erceg, C. J., and A. T. Levin. 2014. "Labor Force Participation and Monetary Policy in the Wake of the Great Recession." *Journal of Money, Credit and Banking* 46 (S2): 3–49.
- Favero, C. A., A. E. Gozluklu, and H. Yang. 2016. "Demographics and the Behavior of Interest Rates." *IMF Economic Review* 64 (4): 732–76.
- Favero, C. A., A. Melone, and A. Tamoni. 2022. "Monetary Policy and Bond Prices with Drifting Equilibrium Rates." *Journal of Financial and Quantitative Analysis* 59 (2): 626–51.
- Galí, J. 2011. "The Return of the Wage Phillips Curve." *Journal of the European Economic Association* 9 (3): 436–61.
- Geanakoplos, J., M. Magill, and M. Quinzii. 2004. "Demography and the Long-Run Predictability of the Stock Market." *Brookings Papers on Economic Activity* 35 (1): 241–325.
- Gertler, M., and P. Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7 (1): 44–76.
- Gilchrist, S., and E. Zakrajšek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102 (4): 1692–1720.
- 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.
- Hirata, W., T. Maruyama, and T. Mineyama. 2020. "Flattening of the Wage Phillips Curve and Downward Nominal Wage Rigidity: The Japanese Experience in the 2010s." Working Paper No. 20-E-4, Bank of Japan.
- Imam, P. A. 2015. "Shock from Graying: Is the Demographic Shift Weakening Monetary Policy Effectiveness." *International Journal of Finance and Economics* 20 (2): 138–54.
- İmrohoroğlu, S., and S. Kitao. 2012. "Social Security Reforms: Benefit Claiming, Labor Force Participation, and Long-Run Sustainability." *American Economic Journal: Macroeconomics* 4 (3): 96–127.
- Inui, M., N. Sudo, and T. Yamada. 2017. "Effects of Monetary Policy Shocks on Inequality in Japan." Working Paper No. 17-E-3, Bank of Japan.

- Iwasaki, Y., I. Muto, and M. Shintani. 2021. “Missing Wage Inflation? Estimating the Natural Rate of Unemployment in a Nonlinear DSGE Model.” *European Economic Review* 132 (February): Article 103626.
- Jaimovich, N., and H. E. Siu. 2009. “The Young, the Old, and the Restless: Demographics and Business Cycle Volatility.” *American Economic Review* 99 (3): 804–26.
- Jarociński, M., and P. Karadi. 2020. “Deconstructing Monetary Policy Surprises — The Role of Information Shocks.” *American Economic Journal: Macroeconomics* 12 (2): 1–43.
- Kondo, A. 2007. “Does the First Job Really Matter? State Dependency in Employment Status in Japan.” *Journal of the Japanese and International Economies* 21 (3): 379–402.
- Leahy, J. V., and A. Thapar. 2022. “Age Structure and the Impact of Monetary Policy.” *American Economic Journal: Macroeconomics* 14 (4): 136–73.
- Lee, K., and S. Ni. 2002. “On the Dynamic Effects of Oil Price Shocks: A Study Using Industry Level Data.” *Journal of Monetary Economics* 49 (4): 823–52.
- Mertens, K., and M. O. Ravn. 2013. “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States.” *American Economic Review* 103 (4): 1212–47.
- Mojon, B., and X. Ragot. 2019. “Can an Ageing Workforce Explain Low Inflation?” BIS Working Paper No. 776.
- Munakata, K., H. Oi, and Y. Ueno. 2019. “Monetary Policy Framework and Information Effect: Evidence from Japanese Data.” Working Paper.
- Muto, I., and K. Shintani. 2020. “An Empirical Study on the New Keynesian Wage Phillips Curve: Japan and the US.” *B.E. Journal of Macroeconomics* 20 (1): 1–17.
- Nakamura, E., and J. Steinsson. 2018. “High Frequency Identification of Monetary Non-Neutrality: The Information Effect.” *Quarterly Journal of Economics* 133 (3): 1283–1330.
- Nakamura, F. 2019. “Household Income, Portfolio Choice and Heterogeneous Consumption Responses to Monetary Policy Shocks.” Discussion Paper No. 2019-E-19, Institute for Monetary and Economic Studies, Bank of Japan.

- Nakamura, F., N. Sudo, and Y. Sugisaki. 2024. "Assessing Monetary Policy Surprises in Japan by High Frequency Identification." *Journal of the Japanese and International Economies* 71 (March): Article 101300.
- Oshio, T., and S. Inagaki. 2015. "The Direct and Indirect Effects of Initial Job Status on Midlife Psychological Distress in Japan: Evidence from a Mediation Analysis." *Industrial Health* 53 (4): 311–21.
- Owan, H. 2004. "Promotion, Turnover, Earnings, and Firm-Sponsored Training." *Journal of Labor Economics* 22 (4): 955–78.
- Ragan, J. F. 1977. "Minimum Wages and the Youth Labor Market." *Review of Economics and Statistics* 59 (2): 129–36.
- Ríos-Rull, J.-V. 2001. "Population Changes and Capital Accumulation: The Aging of the Baby Boom." *B.E. Journal of Macroeconomics* 1 (1): 1–48.
- Romer, C. D., and D. H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94 (4): 1055–84.
- Shimer, R. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95 (1): 25–49.
- Stock, J. H., and M. W. Watson. 2010. "Research Memorandum." https://www.princeton.edu/~mwatson/mgdp_gdi/Monthly_GDP_GDI_Sept20.pdf.
- . 2012. "Disentangling the Channels of the 2007-2009 Recession." *Brookings Papers on Economic Activity* 42 (1): 81–135.
- Stock, J. H., J. H. Wright, and M. Yogo. 2002. "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments." *Journal of Business and Economic Statistics* 20 (4): 518–29.
- Wong, A. 2019. "Refinancing and the Transmission of Monetary Policy to Consumption." Working Paper.
- Yellen, J. L. 2014. "Labor Market Dynamics and Monetary Policy." Speech at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, August 22.
- Zens, G., M. Böck, and T. O. Zörner. 2020. "The Heterogeneous Impact of Monetary Policy on the US Labor Market." *Journal of Economic Dynamics and Control* 119 (October): 1–22.

Monetary Policy Below Zero: An Empirical Investigation of the Reversal Rate*

Zuzana Fungáčová,^a Eeva Kerola,^b and Olli-Matti Laine^b

^aBank of Finland Institute for Emerging Economies (BOFIT)

^bBank of Finland

This study considers the transmission of ECB monetary policy measures to bank corporate lending rates of different maturities from 2010 to 2020. Overall, the transmission of short-term policy rates to lending rates appears to have become weaker when below zero. We observe some signs of the reversal rate during the 2014–20 period, but the evidence is stronger as negative rates become more persistent during the low-for-long period starting in 2016. The emergence of the reversal rate is more pronounced for banks more exposed to negative rates and loans of longer maturities. Unconventional monetary policy measures seem to have mitigated these contractionary effects.

JEL Codes: E52, E58, G21.

1. Introduction

Central banks shifted to accommodative monetary policies in the years following the Global Financial Crisis (GFC) of 2008. A few ventured beyond keeping policy rates close to zero, transitioning into negative territory. On June 5, 2014, the European Central Bank (ECB) became the first major central bank to set policy rates below

*We are grateful to Jin Cao, Jan Fidrmuc, Esa Jokivuolle, Juha Kilponen, Camelia Minoiu, Jaakko Nelimarkka, Salima Ouerk, Aino Silvo, Laurent Weill, and Jurica Zrnc for their useful comments. We are also grateful for the insights from the participants at the Slovak Economic Association Meeting in Bratislava in September 2022, Bank of Finland seminars in 2023, the Research Workshop of the ECB Task Force on Banking Analysis for Monetary Policy in September 2023, and the Dubrovnik Economic Conference in June 2024. This paper represents the views of the authors and not necessarily those of the Bank of Finland. Author emails: zuzana.fungacova@bof.fi (corresponding author), eeva.kerola@bof.fi, olli-matti.laine@bof.fi.

zero, launching an era of negative rates that eventually lasted over eight years. Thus, it is hardly surprising that negative rates have become a subject of extensive study in recent years (Heider, Saidi, and Schepens 2021; Balloch, Koby, and Ulate 2022).

Nevertheless, the large body of theoretical and empirical literature has yet to clarify the overall impact of negative policy rates on banks (Balloch, Koby, and Ulate 2022). There are papers that find no evidence that monetary policy becomes ineffective when rates are negative (Erikson and Vestin 2019; Albertazzi, Nobili, and Signoretti 2021; Altavilla et al. 2022), whereas Wang et al. (2022) show that transmission can already be impaired when the policy rate falls below 0.9 percent. However, some argue that prolonged negative interest rates prompt banks to change their practices, so any observed immediate impact of negative rates, on which the literature has mostly focused so far, may differ from the medium- to long-term impacts (Balloch, Koby, and Ulate 2022). Banks can benefit in the short run from capital gains when negative rates are implemented, helping them to withstand negative rates on the reserves they hold at the central bank. These capital gains are insufficient for the long haul, however, making it hard for banks to maintain profitability. Indeed, Claessens, Coleman, and Donnelly (2018) find that for each additional year of low rates (even if the policy rates stay unchanged), margins and profitability of banks erode by several basis points. Ulate (2021), Abadi, Brunnermeier, and Koby (2023), and Eggertsson et al. (2023) show that when the policy rates are low enough, further rate cuts may become contractionary for bank lending. This reversal interest rate is more likely to emerge if rates are kept low-for-long (Balloch, Koby, and Ulate 2022; Abadi, Brunnermeier, and Koby 2023).

In this paper, we study how the transmission of conventional and unconventional monetary policy to corporate lending rates changes as policy rates go below zero and whether there is empirical evidence of a reversal rate in the euro area. More specifically, we aim to answer the following questions: First, did the pass-through of the ECB's monetary policy to corporate lending rates change when the policy rates became negative and especially when the negative interest rates became more persistent during the low-for-long period? Second, how does the transmission of different unconventional monetary policy measures work in this environment?

These questions are particularly important if there is a higher probability of interest rates staying at lower levels and we face the zero lower bound more often (Holston, Laubach, and Williams 2017; Kiley and Roberts 2017). It is therefore essential to understand how monetary policy affects financial conditions in this environment. The transmission of monetary policy might change as policy rate cuts may even turn contractionary (Abadi, Brunnermeier, and Koby 2023), and unconventional monetary policy measures might be necessary to support the economy.

Our results confirm that when short-term policy rates are lowered below zero, transmission becomes weaker or even reverses. As negative rates become more persistent, we find stronger evidence of the reversal rate—with further rate cuts during the period of low-for-long, banks no longer lower their corporate lending rates but instead raise them. This finding is more pronounced for banks that are more exposed to negative policy rates, i.e., those that are reluctant to lower their own retail deposit rates below zero or those with large amounts of negative interest-bearing central bank deposits. The emergence of the reversal rate also varies between loan maturities, with the inverse effect being most distinct at different time horizons for short-, medium-, and long-term loans. The results indicate that this inverse effect is overall more pronounced for loans with longer maturities. Even if the transmission of the short-term policy rate, a conventional policy measure, to bank lending rates is reversed below zero, unconventional monetary policy measures and targeted longer-term refinancing operations (TLTROs) in particular help to mitigate the pass-through by lowering bank funding costs.

We make several contributions to the existing literature. First, we provide empirical evidence in favor of the theoretically reasoned reversal rate, i.e., the theoretical lower bound for policy-rate lowering, below which cuts become contractionary (Ulate 2021; Abadi, Brunnermeier, and Koby 2023). Moreover, in line with Abadi, Brunnermeier, and Koby (2023), evidence for the reversal rate is stronger after the ECB promised to keep interest rates at negative levels for a prolonged period. Our paper is the first to present empirical evidence of the reversal rate in the euro area. So far, evidence of a reversed transmission has been found for Sweden (Eggertsson et al. 2023) and Switzerland (Basten and Mariathasan 2023). For the euro area, existing literature provides mixed results. Some papers

find that monetary policy is less effective (but not reversed) when interest rates reach very low or negative levels (Borio and Gambacorta 2017; Molyneux et al. 2019; Kwan, Ulate, and Voutilainen 2025). Others, like Erikson and Vestin (2019), Albertazzi, Nobili, and Signoretti (2021), and Altavilla et al. (2022), find no evidence that monetary policy becomes ineffective when rates are negative. Bottero et al. (2022) even show how negative interest rates have expansionary effects, as more liquid banks increase corporate lending more than other banks.

Indeed, our second contribution is to provide an explanation for these mixed results. As suggested by theoretical reasoning, we show that the detection of the reversal rate depends both on bank heterogeneity and loan maturities. Most of the existing literature focuses predominantly on average effects, and thus reversed effects for certain types of banks or loans might have remained undiscovered.

Our third contribution concerns the evidence on how changes in different monetary policy tools transmit to corporate lending rates. Our paper provides clear-cut evidence on how the pass-through of different monetary policy measures to bank lending rates of different maturities has changed over time. The effects of unconventional monetary policy have been studied intensively in recent years (for asset purchases, see Krishnamurthy and Vissing-Jorgensen 2011; for targeted and nontargeted longer-term refinancing operations, see Andrade et al. 2019; Benetton and Fantino 2021; Laine 2021); yet the literature dealing with time-varying effects above and below zero is scarce. Instead, the focus in the low-rate environment has so far been on the effectiveness of short-term policy rates (e.g., Borio and Gambacorta 2017; Claessens, Coleman, and Donnelly 2018; Kwan, Ulate, and Voutilainen 2025), or on unconventional measures but not distinguishing between different kinds of tools (Boeckx, de Sola Perea, and Peersman 2020; Albertazzi, Nobili, and Signoretti 2021). We look separately at the time-varying effects of short-term policy rates, quantitative easing (QE), and (T)LTROs.

As we aim to investigate the existence of the reversal rate, we focus on bank lending rates. Developments in lending rates are found to be mostly driven by supply shocks, whereas fluctuations in lending volumes are largely explained by demand (Altavilla, Boucinha, and Bouscasse 2022). We exploit a detailed bank-level data set of

corporate lending rates at the monthly frequency on 137 individual banks from 13 euro-area countries covering the period from January 2010 to December 2020. Unlike most empirical papers that rely on shorter data samples,¹ our data enable us to properly study the low-for-long period. For our purposes, the low-for-long period begins in 2016, when ECB forward guidance signaled that policy rates would remain in negative territory “for a prolonged time.” Moreover, we utilize the detailed information on different loan maturities. Besides information on bank balance sheets, loan and deposit interest rates, central bank deposits, and longer-term borrowing of banks, we employ data from the responses of individual banks to the ECB’s Bank Lending Survey to control for bank-level loan demand. We isolate monetary policy shocks by utilizing rate changes around the ECB’s monetary policy decisions from the Euro Area Monetary Policy Event-Study Database (EA-MPD; see Altavilla et al. 2019) and bank bond yield changes on (T)LTRO announcement days. These shock proxies allow us to disentangle various types of monetary policy tools: short-term policy rates, targeted long-term refinancing operations, and quantitative easing.

The article has the following structure. Section 2 describes the main aspects of the monetary policy conducted by the ECB during the observation period 2010 to 2020. Section 3 presents our data and various measures of monetary policy. Section 4 provides the methodology, and Section 5 presents the results. Section 6 concludes with a few policy insights.

2. Monetary Policy in the Euro Area


This section presents the main aspects of the ECB’s monetary policy during the period from 2010 to 2020.² It also discusses how we define the low-for-long period.

Policy rate-setting is the ECB’s primary monetary policy tool (Figure 1). Its main refinancing operations (MROs) allow banks to borrow from it on a weekly basis at predetermined rates. Liquidity sharing on the interbank market stalled with the onset of the

¹See Table 4 in Balloch, Koby, and Ulate (2022) for details concerning the data used in existing empirical studies.

²For a more comprehensive description, see Rostagno et al. (2021).

Figure 1. ECB Policy Rates and EONIA (2010–20)




Source: ECB.

GFC, and in October 2008, the ECB began carrying out its MROs as fixed-rate tenders with full allotment. This shift led to a buildup of liquidity inside the banking sector. Because of this excess liquidity, the shortest money market rates (e.g., the Euro Overnight Index Average, or EONIA) began to track the ECB's deposit facility rate rather than the MRO.

Following the ECB's initial cut of its deposit facility rate to – 0.10 percent in June 2014, key interest rates were lowered a total of five times. The final cut, in September 2019, brought the deposit facility rate to –0.50 percent. When the deposit facility rate is negative, banks have less to withdraw from their central bank accounts than their previous day's deposit. Individual banks can reduce their own excess liquidity by lending it out to other banks or purchasing assets, but as liquidity is always passed on from one bank to another, the banking system as a whole cannot shed its total excess liquidity. This burden was reduced in October 2019 with the adoption of a two-tier system for reserve remunerations.

The ECB also contributed to lowering banks' funding costs by lending funds on highly favorable terms under both nontargeted and targeted longer-term refinancing operations (Figure 2). Two LTROs with a maturity of three years were agreed upon in December 2011. The operations became targeted (i.e., TLTROs) in 2014 as the aim

Figure 2. ECB's Longer-Term Refinancing Operations and Asset Purchases under the Asset Purchase Programme (APP) and Pandemic Emergency Purchase Programme (PEPP)



Source: ECB.

was refined such that banks would use the new liquidity to increase lending to nonfinancial corporations and households for consumption. The cost of TLTROs was bank-specific and became cheaper the more the recipient bank increased its lending to the private sector. The first series of TLTROs (TLTRO I) was announced in June 2014, followed by TLTRO II in March 2016, and TLTRO III in March 2019.

Additional excess liquidity was also created by the ECB's Expanded Asset Purchase Programme (APP). The Eurosystem conducted net purchases of securities under several asset purchase programs (the Public Sector Purchase Programme (PSPP), the Third Covered Bond Purchase Programme (CBPP3), the Corporate Sector Purchase Programme (CSPP), and the Asset-Backed Securities Purchase Programme (ABSPP)) at varying monthly purchase paces between October 2014 and December 2018. APP net purchases stalled between January and October 2019 but were restarted in November 2019 and kept running at a pace of 20 billion euros a month until the end of our observation period in December 2020. In March 2020, as the first wave of the COVID-19 pandemic was hitting Europe, the ECB launched its Pandemic Emergency Purchase Programme (PEPP). It made additional purchases in all

categories eligible under the APP (plus some extensions such as Greek government securities).

Our analysis focuses on the changes in the transmission of monetary policy through banks (i) after policy interest rates fall into negative territory, and (ii) when economic agents realize that negative rates are here to stay, i.e., when the economy enters the low-for-long period.


The first deposit facility rate cut below zero in June 2014 marked the beginning of the negative rates era that lasted until the end of our sample. However, in our estimations, we extend the period of negative rates to January 2014 in order to make sure that any surprises related to the transition into negative territory are also included in the subsample.

Defining the start date for the low-for-long period is more challenging. The ECB lowered its deposit facility rate a third time in December 2015, from -0.20 to -0.30 percent, and again to -0.40 percent in March 2016. On January 21, 2016, ECB President Mario Draghi announced a shift in the ECB's forward guidance with the following statement after the monthly monetary policy meeting:

Based on our regular economic and monetary analyses, and after the recalibration of our monetary policy measures last month, we decided to keep the key ECB interest rates unchanged, and we expect them to remain at present or lower levels for an extended period of time. (ECB 2016)

The change in the expected interest rate path was significant after these policy changes. Panel A of Figure 3 depicts the expected path (forward curve) of the euro-area short-term rate and its distributions at two different points in time: the second rate cut in September 2014 (pink distribution) and the fourth rate cut in March 2016 (blue distribution). These expectations are derived from the EURIBOR (euro-area interbank offered rate) future option market assuming risk-neutrality following the method of Shimko (1993). Note that even after the policy rate was lowered to -0.20 percent, the scenario in which interest rates would return to positive territory over the course of one year was within 50 percent confidence bands. However, entering 2016, this perception had changed. Looking at the probability distribution in March 2016, markets clearly started to understand that negative interest rates were a persistent phenomenon.

Figure 3. Expected Euro-Area Short-Term Rate Paths and Their Option-Implied Probability Distributions (Panel A). Share of Banks Having at Least One Retail Deposit Account at Negative Rate (Panel B)



Source: ECB; authors' calculations.

The change in perception of the persistence of negative interest rates is also visible when examining the number of banks in our data set that started to charge negative interest rates on their retail deposits. As shown in panel B of Figure 3, the share of banks with at least one retail deposit account (including both household and corporate, overnight, and other maturities) bearing negative interest rates increased considerably from the beginning of 2016 to around 20 percent by the end of 2017 and 30 percent in 2020. Banks that started to charge negative interest rates on their retail deposits were on average not different from their nonnegative deposit rate counterparts with respect to their ex ante (or ex post) level of net interest margin. In addition, charging negative deposit rates does not correlate with the amount of variable interest lending nor the initial level of either lending or deposit rates. Out of the 13 countries in our sample, there were in total 9 countries where we observed banks charging negative deposit rates. Based on the reasoning above, we place the beginning of the low-for-long period in January 2016, and it lasted until the end of our observation period in December 2020.

3. Data

This section presents the data used and defines our monetary policy measures.

3.1 *Data Description*

We employ different sources of data to create our unique data set. First, our dependent variables, i.e., the lending rates for new loans, are drawn from the ECB's confidential individual MFI interest rate statistics (IMIR). We have bank-specific interest rate observations at a monthly frequency on new domestic nonfinancial corporation³ (henceforth corporate) loans of different maturities. We use four dependent variables for our main estimations. The average lending rate for corporate loans is a weighted average of lending rates for all new corporate loans of different maturities granted by a bank in each month. We also look at the lending rates for corporate loans separately for maturities of up to one year, one to five years, and over five years.

Our main explanatory variables are the measures of monetary policy. We use a local projection method to disentangle the effects of specific monetary policy measures: policy interest rate, (T)LTROs, and QE. To that end, we use the EA-MPD and our own calculations. Monetary policy measures used are discussed more thoroughly in Section 3.2.

To form our bank-specific variables, capitalization (ratio of equity to total assets), liquidity (ratio of liquid assets to total assets), and bank size (log of total assets), we predominantly rely on monthly bank-specific balance sheet information provided by the ECB (individual MFI balance sheet item data set, IBSI); the exception being liquidity, where we use annual data from BankFocus. In addition, we use monthly observations on each bank's borrowing from the central bank under targeted and nontargeted longer-term operations, as well as each bank's monthly reserves in central bank deposit accounts. The data are from two confidential ECB data sets (bank-level borrowing and repayments from different longer-term refinancing operations and bank-level current account deposits).

With access to the ECB's individual Bank Lending Survey data with quarterly survey responses of individual banks, we have the

³We focus on corporate loans to keep our analysis concise. Nevertheless, the examination of housing loans yields similar results. These results are available upon request.

possibility to control directly for loan demand encountered by each bank. Our question of interest is worded as follows:

Over the past three months (apart from normal seasonal fluctuations), how has the demand for loans or credit lines to enterprises changed at your bank? Please refer to the financing need of enterprises independent of whether this need will result in a loan or not.

Respondents could choose from “(i) decreased considerably, (ii) decreased somewhat, (iii) remained basically unchanged, (iv) increased somewhat, or (v) increased considerably.” Based on the responses, we construct two dummies for loan demand: one for increasing demand (equal to one if the bank responded with (iv) or (v)), and one for declining demand (equal to one if the bank responded with either (i) or (ii)). Each estimation includes both dummies.

To control for macroeconomic developments affecting all banks in one country at the same time, we include country-specific year-on-year growth in industrial production and the harmonized unemployment rate, both from Eurostat at a monthly frequency. In addition, to take into account euro-area-wide macroeconomic developments and uncertainty, we further include normalized total returns of the euro-area STOXX index, as well as the first principal component of short-, medium-, and long-term expectations of euro-area GDP growth, inflation, and unemployment from the ECB Survey of Professional Forecasters (SPF) at monthly frequency.

Altogether, our data set is an unbalanced panel covering 137 individual euro-area banks from 13 countries (Austria, Belgium, Estonia, Finland, France, Germany, Ireland, Italy, Lithuania, Luxembourg, Portugal, Slovakia, and Spain) for the period January 2010–December 2020. The smaller coverage of the individual Bank Lending Survey (iBLS) ultimately limits our sample. Even if the number may seem small at first, banks participating in the BLS provide a representative sample of euro-area banks and take into account the characteristics of the respective national banking structures (see also Figure A.1 in the appendix). Definitions of all variables used are presented in Table 1. Descriptive statistics for the full time period and two different subperiods (period of negative policy rates and period of low-for-long) can be found in Table 2. We

Table 1. Definitions of Variables

Variable	Definition	Source
NFC Lending Rate	NFC loan rate weighted average (weighted average calculated using the amount of new loans provided every month)	ECB IMIR
NFC Lending Rate, up to One Year	NFC loan rate for new loans up to one year of maturity	ECB IMIR
NFC Lending Rate, One to Five Years	NFC loan rate for new loans one to five years of maturity	ECB IMIR
NFC Lending Rate, over Five Years	NFC loan rate for new loans over five years of maturity	ECB IMIR
Liquidity Capitalization Size	Liquid assets to total assets ratio Equity to total assets ratio Log of bank total assets	Bank Focus ECB IBSI ECB IBSI
Central Bank Deposits	Banks' monthly reserves in central bank deposit accounts	ECB confidential database
Central Bank Operations	Banks' borrowings from the central bank's targeted and non-targeted longer-term operations	ECB confidential database
Industrial Production	Country-specific industrial production year-on-year growth	Eurostat
Unemployment	Country-specific harmonized unemployment rate	Eurostat
Stock Return	Euro area, STOXX index, total return, close (in EUR), normalized	STOXX, Macrobond
Exp. Euro-Area Inflation	Expected rate of inflation in the euro area, 1st principal component of short-, medium-, and long-term expectations	ECB Survey of Professional Forecasters
Exp. Euro-Area Unemployment	Expected rate of unemployment in the euro area, 1st principal component of short-, medium-, and long-term expectations	ECB Survey of Professional Forecasters
Exp. Euro-Area GDP	Expected rate of GDP in the euro area, 1st principal component of short-, medium-, and long-term expectations	ECB Survey of Professional Forecasters
Krippner	Euro-area shadow rate by Krippner (2015)	Macrobond
Demand Increase	Dummy for increasing NFC loan demand (equals one if a bank responded either 4 or 5)	ECB individual Bank Lending Survey
Demand Decrease	Dummy for decreasing NFC loan demand (equals one if a bank responded either 1 or 2)	ECB individual Bank Lending Survey

Table 2. Summary Statistics

Variable	Full Time Span (2010–20)			Negative Policy Rates (2014–20)			Period of Low-for-Long (2016–20)		
	# of Obs	Mean	Std. Dev	# of Obs	Mean	Std. Dev	# of Obs	Mean	Std. Dev
NFC Lending Rate	11,557	2.366	1.111	7,821	2.037	0.963	5,495	1.853	0.855
NFC Lending Rate, up to One Year	11,426	2.312	1.102	7,756	2.013	0.961	5,464	1.841	0.854
NFC Lending Rate, One to Five Years	10,407	2.806	1.637	7,060	2.223	1.348	5,008	1.942	1.202
NFC Lending Rate, over Five Years	9,880	2.995	1.614	6,813	2.418	1.314	4,910	2.138	1.154
Liquidity	13,020	0.187	0.147	8,616	0.193	0.150	6,012	0.207	0.150
Capitalization	12,519	0.096	0.049	8,463	0.101	0.049	5,957	0.100	0.047
Size	12,519	10.780	1.487	8,463	10.732	1.509	5,957	10.759	1.507
Central Bank Deposits	13,020	0.013	0.031	8,616	0.017	0.033	6,012	0.020	0.036
Central Bank Operations	13,020	0.190	4.618	8,616	0.286	5.675	6,012	0.404	6.790
Industrial Production	13,020	1.206	7.267	8,616	1.022	7.798	6,012	0.254	7.428
Unemployment	13,020	9.176	5.009	8,616	8.481	4.495	6,012	7.769	3.908
Stock Return	13,020	0.037	0.971	8,616	0.634	0.561	6,012	0.862	0.486
Exp. Euro-Area Inflation	13,020	-0.011	1.630	8,616	-0.895	1.180	6,012	-0.698	1.256
Exp. Euro-Area Unemployment	13,020	0.044	1.736	8,616	-0.507	1.713	6,012	-1.436	1.110
Exp. Euro-Area GDP	13,020	-0.023	1.188	8,616	0.469	1.152	6,012	0.755	1.256
Krippner	13,020	-1.533	1.149	8,616	-2.162	0.666	6,012	-2.289	0.726
Demand Increase	13,020	0.175	0.380	8,616	0.213	0.409	6,012	0.219	0.414
Demand Decrease	13,020	0.158	0.365	8,616	0.121	0.326	6,012	0.112	0.316

exclude the extreme values by dropping 1 percent of observations from each side of the distribution of the dependent variables.

3.2 Measuring Monetary Policy

As discussed in Section 2, the ECB employed a number of monetary policy instruments during the 11 years covered by our data set. Our focus is to determine whether the pass-through of monetary policy to bank corporate lending rates changed after policy interest rates entered into negative territory and further when banks entered the low-for-long period in January 2016.

A standard strategy in the recent literature is to use monetary policy shocks derived from external VAR models (e.g., Boeckx, de Sola Perea, and Peersman 2020 utilize shocks from Boeckx, Dossche, and Peersman 2017) or event studies (e.g., Albertazzi, Nobili, and Signoretti 2021; Ampudia, Ehrmann, and Strasser 2023; Kho 2025). The monetary policy shocks and surprises represent exogenous variation in monetary policy, and they also allow us to disentangle different types of monetary policy.⁴

We employ the local projections method and three different monetary policy surprises to disentangle the impacts of short-term policy rates, credit easing policies ((T)LTROs), and large-scale asset purchases (QE).

We take the one-week overnight index swap (OIS) surprise around the ECB's monetary policy decisions using the monetary event window (that is, the change in the median quote from the window 13:25–13:35 before the press release to the median quote in the window 15:40–15:50 after the press conference) to measure the impact of conventional monetary policy (the short-term interest rate) and the 10-year OIS surprise to measure the impact of large-scale asset purchases or QE. Both of these surprises are derived from the euro-area event-study database produced in accordance with Altavilla et al. (2019). When studying QE using the 10-year

⁴Alternatively, we could measure monetary policy using an endogenous (shadow) policy rate and control for observed variables in the central bank reaction function as discussed by Bluedorn, Bowdler, and Koch (2017). This approach was utilized in the earlier version of this article. The results based on this alternative strategy are in line with the results presented in the latter sections.

OIS surprise, we use the one-week surprise as a contemporaneous control variable because long-term interest rates are affected by changes in the short-term rates. To measure credit easing policies, we use bank bond yield changes around the ECB announcements involving (T)LTROs when controlling for other changes in monetary policy using the one-week and 10-year surprises. We use the same days as in Altavilla et al. (2023, Appendix B): May 8, 2014, June 5, 2014, July 3, 2014, July 29, 2014, January 22, 2015, March 10, 2016, and May 3, 2016. In addition, we add the following days: December 8, 2011 (three-year LTRO announcement), February 28, 2012 (eligibility of Greek bonds used as collateral in monetary policy operations), March 7, 2019 (TLTRO III announcement), June 6, 2019 (details about the operations), July 29, 2019 (legal act published), and September 9, 2019 (TLTRO rate reduced). We use a single series of average yield to maturity of euro-area banks (senior unsecured bonds), and based on this single series, we calculate the daily change. The data are from Bloomberg.


Figure 4 presents the time series of these three surprises for 2010–2019. We do not use data concerning the shock proxies after 2019 due to extreme values related to the COVID-19 pandemic.

4. Methodology

We utilize a local projections model as formalized by Jordà (2005), although the idea of estimating regression coefficients of a model where the left-hand-side variable is several periods ahead of the right-hand-side variable is of course much older (e.g., Cox 1961). This approach enables us to estimate the effects of policy shocks at different horizons and thereby tease out the timing and dynamics of these effects. Furthermore, high-frequency monetary policy surprises provide us with an identification strategy for causal effects. In line with the recent literature (Boeckx, de Sola Perea, and Peersman 2020; Altavilla et al. 2022; Altavilla et al. 2023; Kho 2025), we estimate a model that applies local projection methods to bank-level data:

$$r_{i,t+h}^L = a_{i,h} + \rho_h(L)Y_{t-1} + \gamma_h(L)X_{i,t-1} + \theta_h Shock_t + e_{i,t+h} \quad (1)$$

**Figure 4. Monetary Policy Shocks
Used in Local Projections**



Source: EA-MPD; Bloomberg; authors' calculations.

for different horizons h . In the equation, $r_{i,t}^L$ is the bank lending rate for firms at different maturities, L , charged by bank i at time t ; $a_{i,h}$ are bank-specific fixed effects (for different h); Y_t is the vector of macroeconomic control variables; and $X_{i,t}$ is the vector of bank-specific control variables. $Shock_t$ stands for proxy variables for monetary policy shocks (we use three shocks: a policy rate shock, a (T)LTRO shock, and a QE shock), defined in more detail in Section 3.2. The majority of the recent literature (Boeckx, de Sola Perea, and Peersman 2020; Ampudia, Ehrmann, and Strasser 2023; Altavilla et al. 2023; Kho 2025) includes the shock proxy directly in the equation as we do. Some papers, such as Altavilla et al. (2022), use shock proxies as an instrumental variable for the endogenous policy rate. We assess the robustness of our main results regarding conventional monetary policy using this alternative approach. Following the application of local projections by Boeckx, de Sola Perea, and Peersman (2020) to similar data, we assume the number of lags in the baseline specification to be three. The regressions are estimated using OLS. In all the figures, we show 90 percent confidence intervals for θ_h . The

confidence intervals are calculated based on a nonparametric robust covariance matrix estimator à la Driscoll and Kraay (1998).

We control for expectations about the macroeconomic environment at the euro-area level by including GDP growth, inflation, and unemployment forecasts from the Survey of Professional Forecasters (SPF), and the current environment by including normalized Eurostoxx stock returns, as well as country-specific industrial production growth and unemployment rates to control for country-specific macroeconomic developments affecting all banks inside a country. Bank-specific controls include liquidity, capitalization, and bank size. We account for corporate loan demand using bank-specific responses from the ECB's individual Bank Lending Survey (iBLS) on how demand for loans or credit lines to enterprises has changed. We use two dummy variables in each estimation: an increasing demand dummy and a declining demand dummy. In addition, we control for the stance of monetary policy prior to the shock using the EONIA and Krippner shadow rates.

Because our proxy variables for monetary policy shocks can be assumed to be exogenous, control variables are not necessary for identification. Nevertheless, adding control variables to the regression reduces the confidence intervals and helps produce more accurate estimates. In addition, adding macro controls to the model removes a potential concern that, for example, more stimulating surprises might occur when the economic situation is bad. The control variables are lagged because we do not want to control out a potential effect of monetary policy on lending rates via immediate effects on the variables used as controls (e.g., Jordà 2005; Plagborg-Møller and Wolf 2021). In some cases, as explained in Sections 5.2 and 5.3, we add one or more contemporaneous policy surprises as control variables to disentangle different types of monetary policies from each other. For example, when studying asset purchases via a surprise change in the long-term rate, we add a one-week rate surprise as a control variable because the long-term rate may be affected by contemporaneous short-term rate surprises.

5. Results

We aim to study how the transmission of conventional and unconventional monetary policy to corporate lending rates changes as

policy rates go below zero. To do so, we analyze whether there are differences in the effects of specific monetary policy tools used by the ECB. Such differences are crucial from a policy standpoint. Section 5.1. starts with the conventional monetary policy tool, the short-term policy rate. Section 5.2. moves on to longer-term refinancing operations, and Section 5.3. looks at asset purchase programs.

5.1 Policy Interest Rate

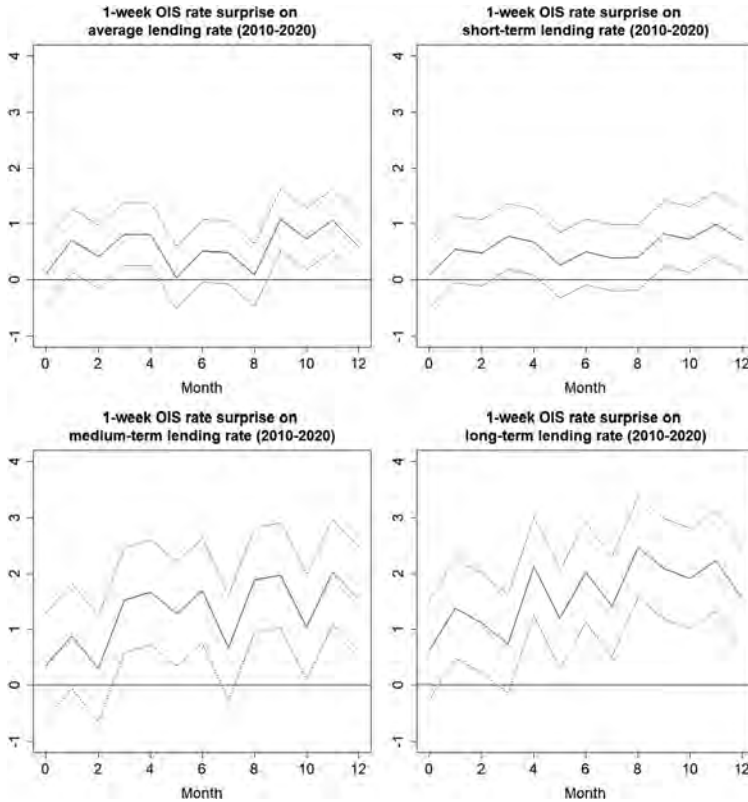
As discussed in Section 4, the short-term policy rate is proxied by a one-week OIS rate shock.⁵ We use full calendar years and overlapping samples: (i) the full time span (2010–20), (ii) the period of negative rates (2014–20), and (iii) the period of low-for-long (2016–20).

Figure 5 shows the local projection results for 12-month horizons for the average corporate loan rate and three different maturities for the full time span (2010–20). The peak effects on the average rate are reported in Table A.2. The results suggest that (unexpected) short-term rate changes transmit to bank lending rates approximately one-to-one. When it comes to the average lending rate (upper left panel of Figure 5), the peak effect materializes nine months after the shock (at horizon 9). The point estimate is about 1, which means that, e.g. a negative short-term rate surprise of 25 basis points induces a 25 basis point decrease in the average lending rate. There are some differences with the timing of the peak effect when looking at loans of different maturities. Transmission seems to be stronger overall for loans of longer maturities.

Once policy rates fall below zero, things change. Panels in Figure 6 suggest that the pass-through of policy rate changes to bank lending rates is greatly impaired below zero, as they no longer have the same impact on these rates. For the average corporate lending rate, the effect is statistically not different from zero. The same applies for lending rates of short- and medium-term loans. In the case of the long-term lending rate, results start providing a hint of the reversal rate, as the effect is negative and statistically significant around six months after the shock. Following a 25

⁵We also use a one-month rate surprise instead of the one-week rate surprise. This does not affect the results.

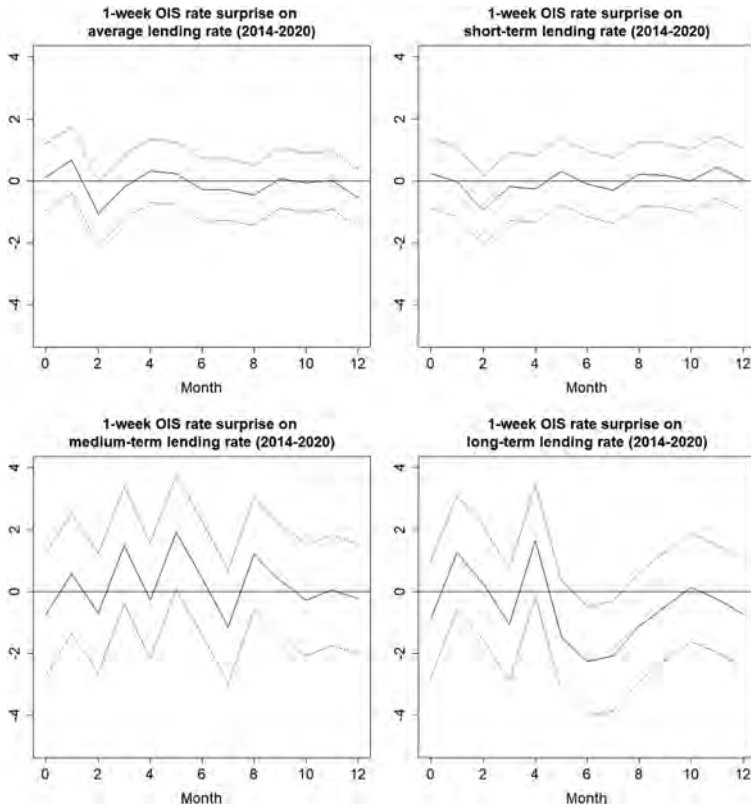
Figure 5. Effect of Short-Term Rate Shock for the Average Corporate Lending Rate and Three Different Maturities, Full Observation Period (January 2010–December 2020)



Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

basis point negative short-term rate surprise, banks increase their long-term lending rates by 50 basis points. Note that in this particular exercise we utilize both positive and negative short-term rate surprises during the period of negative rates. Even though rates were not raised after 2014, market participants at times anticipated more aggressive rate cuts than those that were actually decided

Figure 6. Effect of Short-Term Rate Shock for the Average Corporate Lending Rate and Three Different Maturities, Period of Negative Rates (January 2014–December 2020)

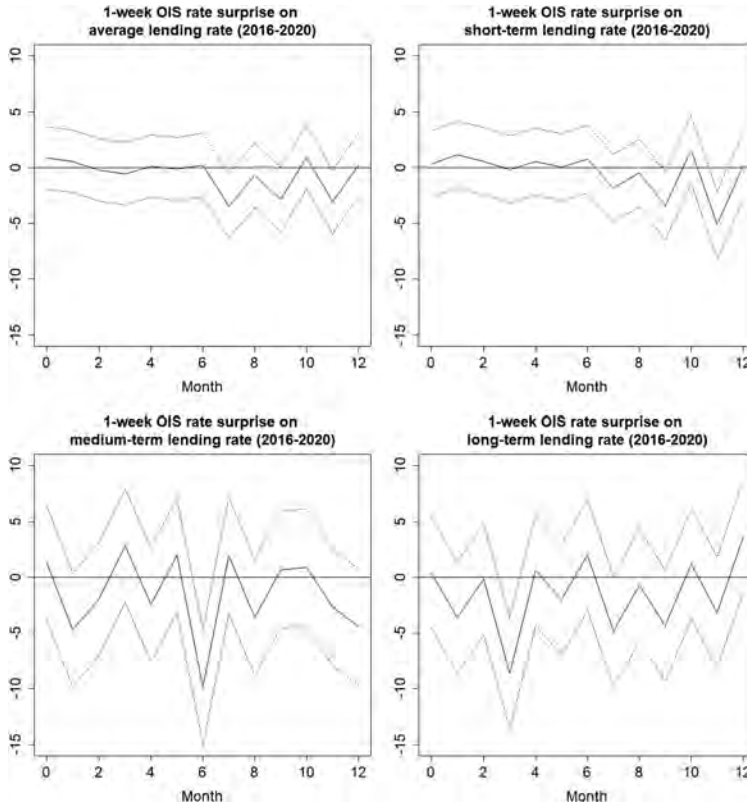


Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

at ECB policy meetings (i.e., contractionary rate surprises). Using only the negative rate surprises for this subsample does not alter the results (see Figure A.2 in the appendix).

As negative rates become a more persistent phenomenon, we again see changes in the estimated coefficients. Figure 7 presents the results for the low-for-long period subsample, which runs from

Figure 7. Effect of Short-Term Rate Shock for the Average Corporate Lending Rate and Three Different Maturities, Low-for-Long Period (January 2016–December 2020)



Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

2016 to 2020. We now see more frequent instances of horizons with negative and statistically significant coefficients, suggesting the existence of the reversal rate, where additional (here unexpected) rate cuts made banks increase their corporate lending rates. The estimated reversal effect for the average lending rate is economically significant. The peak effect at horizon 7 is about -3 , meaning that

a 25 basis point negative short-term rate surprise further into the negative territory would raise the average lending rate by about 75 basis points. The loan maturities again play a role, as the peak effect occurs at different horizons for different maturities. The statistical significance and the absolute size of the coefficients are higher for medium- and long-term loans.

We assess the robustness of our results concerning policy rates in multiple ways. First, Table A.1 in the appendix provides results using the rolling window approach instead of subjectively chosen subsamples. Results confirm that negative and statistically significant coefficients for the average lending rate can only be found for time windows that incorporate at least the year 2016, so that the reversal rate emerges only when looking at the subsamples containing the low-for-long period. We observe positive, statistically significant coefficients for time windows ending in 2015 or earlier. Second, Figure A.3 presents results for the average lending rate accounting for the future values (leads) of expected GDP growth and loan demand dummy variables, as monetary policy could also affect lending rates indirectly via its dynamic effect on loan demand.⁶ This does not change the results. Third, Figure A.4 reports results for the main specification of the average lending rate after replacing the country-specific macro variables in the baseline specification by country and country-year fixed effects to control for any country-specific variation. Here again, the results corroborate our main results. Fourth, although our focus is on lending rates when investigating the existence of the reversal rate (as explained in Section 1), we complement our analysis with regressions for loan amounts (Figure A.5). Using the full sample or the sample prior negative rates, the results suggest that monetary policy easing has a positive effect on loan volumes.

⁶Additionally, instead of a one-week OIS surprise, we use the Target shock constructed as in Altavilla et al. (2019). We use the replication of this shock proxy from Laine and Pihlajamaa (2024) and the poor man's method by Jarociński and Karadi (2020) to clean out the potential information effects. As the correlation between the high-frequency shock proxies and money market rates might differ across subsamples, we use the shock proxy as an instrumental variable for the change in the three-month Euribor rate. For the period of negative rates (2014–20) we find evidence for a reversal rate. Focusing only on the period of low-for-long yields statistically insignificant results. This is not surprising as the number of nonzero values in the shock proxy is very low after 2016 in this specification. These results are available upon request.

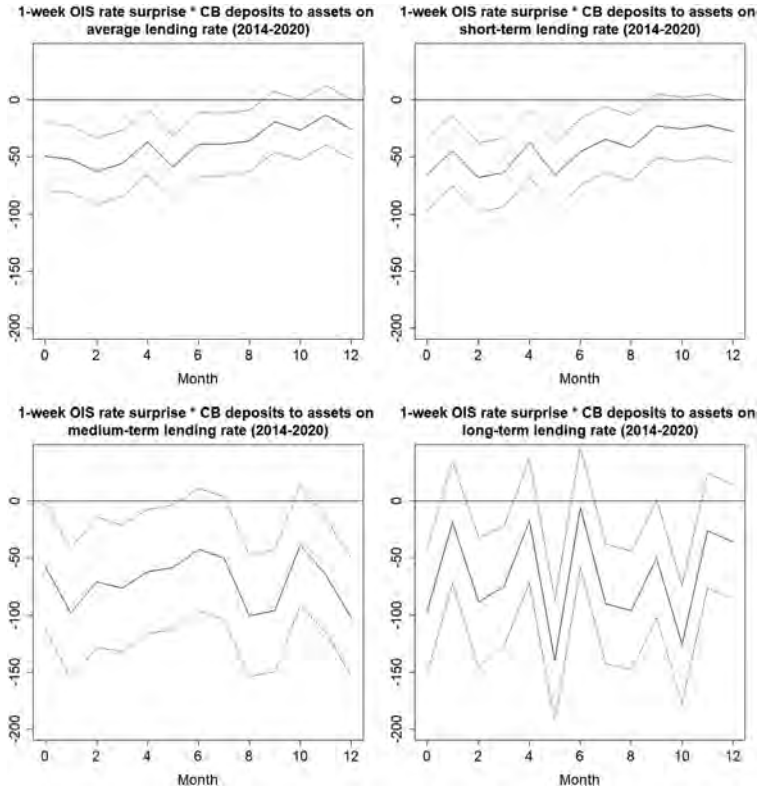
During the period of low-for-long, things change. In the same way as for loan rates, we find evidence that an unexpected policy rate surprise has an inverse effect on loan volumes. The results are in line with the idea of the reversal rate. Nevertheless, we acknowledge that the results may capture equilibrium relationships between policy rate changes and banks' loan pricing decisions, rather than pure supply-side effects, as firms' demand for credit may react immediately to monetary policy. We assess the role of loan demand by controlling for the future values of our loan demand control variable in Figure A.3 in the appendix. The results remain qualitatively the same, but we cannot rule out that our control variables do not capture all the changes in credit demand.

5.1.1 Bank Heterogeneity

To understand our reversal rate results better, we assess if the negative effect is stronger for banks that are more exposed to negative rates. Following Demiralp, Eisenschmidt, and Vlassopoulos (2021), we analyze whether banks with a high share of central bank deposits behave differently from the average bank. We augment Equation (1) to include an interaction term of our shock proxy and the lagged share of central bank deposits on bank i 's balance sheet. Figure 8 presents the results, providing evidence that banks with more central bank deposits (i.e., bearing the additional cost of the negative central bank deposit rate) raise lending rates for their borrowers. At horizon 2, the coefficient estimate for the interaction term is about -60 and the estimate for the coefficient of the one-week OIS surprise is about -0.4 . This means that for a bank with central bank deposits to total assets at 10 percent, a 25 basis point negative short-term rate surprise leads to an increase in the average lending rate of about 160 basis points. In turn, a bank with a central bank deposit ratio of 1 percent raises its average lending rate by 25 basis points. This result corroborates the theoretical reasoning of Ulate (2021).

Bank heterogeneity also matters in how banks are able to keep lowering their interest expenses as policy rates go below zero. As argued by Abadi, Brunnermeier, and Koby (2023) and Eggertsson et al. (2023), the key constraint preventing policy rate pass-through to bank lending rates in a below-zero environment is that some banks cannot lower their retail deposit rates into negative territory. We

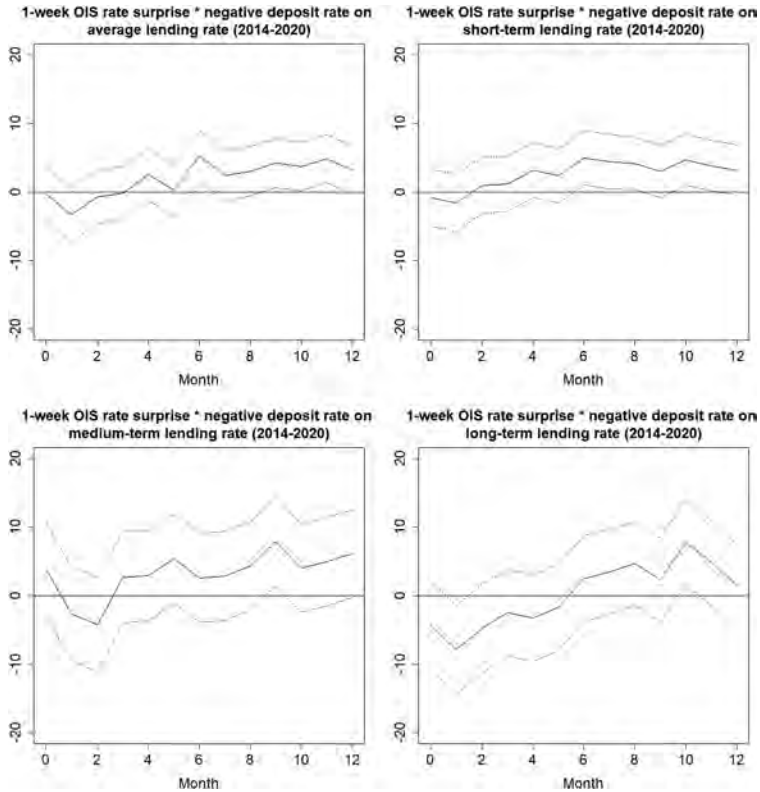
Figure 8. Effect of Short-Term Policy Rate Shock on Banks with High Share of Central Bank Deposits for the Average Corporate Lending Rate and Three Different Maturities, Period of Negative Rates



Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

address this question by augmenting Equation (1) with the interaction term between our short-term policy rate shock and a lagged dummy variable that has a value of 1 if a bank sets at least one of its retail deposit rates to negative. Figure 9 reports the coefficient of this interaction term for different horizons.

Figure 9. Effect of Short-Term Rate Shock on Banks with Negative Retail Deposit Rates for the Average Corporate Lending Rate and Three Different Maturities, Period of Negative Rates



Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

Our results indicate that banks that manage to lower their retail deposit rates below zero are still transmitting changes in interest rates to their borrowers. This is fully in line with the reasoning of Abadi, Brunnermeier, and Koby (2023), and Eggertsson et al. (2023), who argue that those banks that can further decrease their interest expenses can pass on costs from changes in the short-term

policy rate to their customers in the form of lower lending rates. Indeed, these results clearly demonstrate that bank heterogeneity matters for the transmission of policy rates to bank lending and may be a defining factor in whether further rate cuts below zero become contractionary.

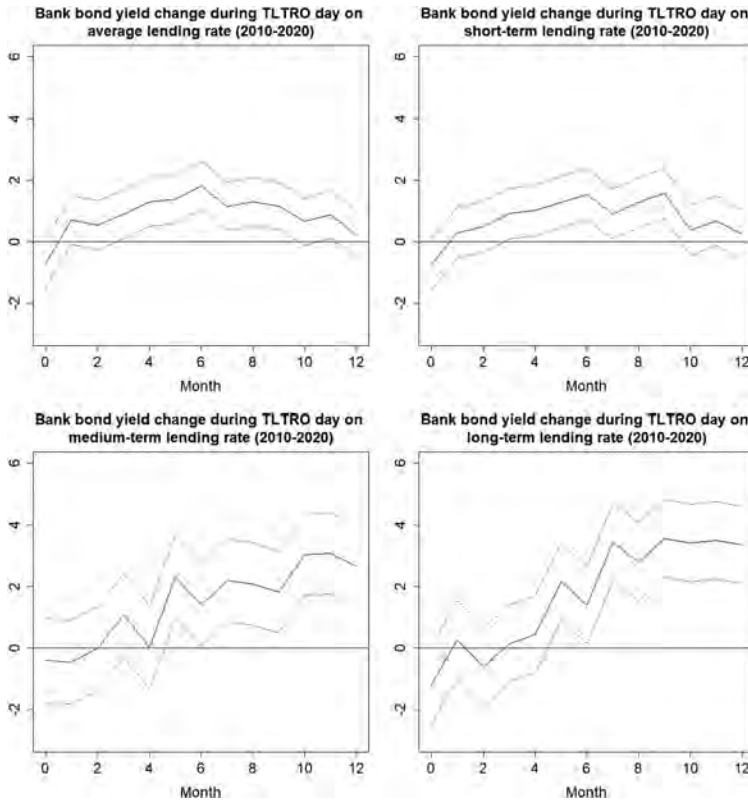
5.2 (T)LTROs

Besides eventually lowering the conventional monetary policy measure below zero, the ECB introduced a variety of different unconventional monetary policy tools. In this subsection, we study the transmission of those instruments to bank lending rates. Our (T)LTRO shock proxy series thus includes two LTRO-related yield movements and all variations after 2014 are related to targeted operations. Equation (1) now includes (T)LTRO-related policy changes proxied by the average bank bond yield change during the days of ECB announcements. In addition, we augment the equation with contemporaneous one-week and ten-year rate surprises to ensure that our (T)LTRO proxy is not contaminated by other monetary policy tools.

The results for the full sample are shown in Figure 10. They suggest that these policies affect lending rates with some lag, as the full effect does not materialize until after a year has passed. Following a 10 basis point decrease in (T)LTRO-induced bank bond yields, we observe a roughly 20 basis point decrease in the average lending rate. As with short-term policy rates, this effect is larger and more pronounced for lending rates of longer maturities.

Figure 11 shows the results for the period of negative rates. The estimated effects in this subsample are much smaller and statistically less significant (although positive) for the average corporate lending rate and the lending rate for short-term loans. For medium and long maturities, positive and statistically significant results remain. The slightly stronger positive effect for long-term lending rates remains during the low-for-long period (Figure 12). Even so, caution is warranted in interpreting these results due to the tiny number of TLTRO announcement days for this last subperiod. In addition, the time-varying effectiveness of TLTROs likely reflects differences in their “calibration:” the first, second, and third series of TLTROs differ when it comes to maturity and incentives.

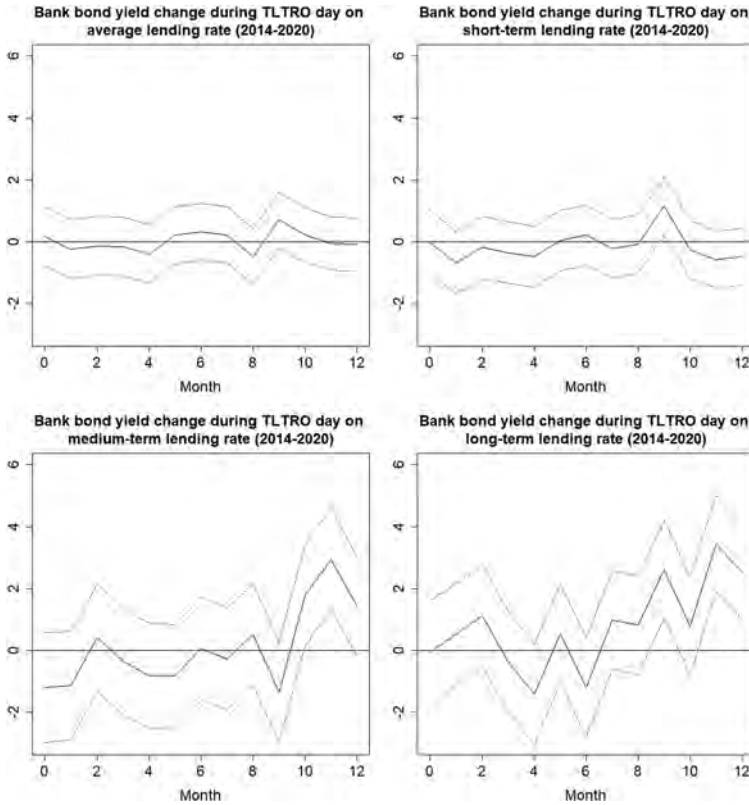
Figure 10. Effect of (T)LTRO Shock for the Average Corporate Lending Rate and Three Different Maturities, Full Observation Period (January 2010–December 2020)



Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

The results above assess the effect on the lending rates of an average bank. However, one may assume that the effects are stronger for banks that participated in these operations and received more financing on favorable terms. This is in line with Gertler and Kiyotaki (2010), where an injection of liquidity from a lender of last resort relieves the bank’s credit constraint, which then further alleviates liquidity shocks to the nonfinancial sector. Our data set includes

Figure 11. Effect of (T)LTRO Shock for the Average Corporate Lending Rate and Three Different Maturities, Period of Negative Rates (January 2014–December 2020)

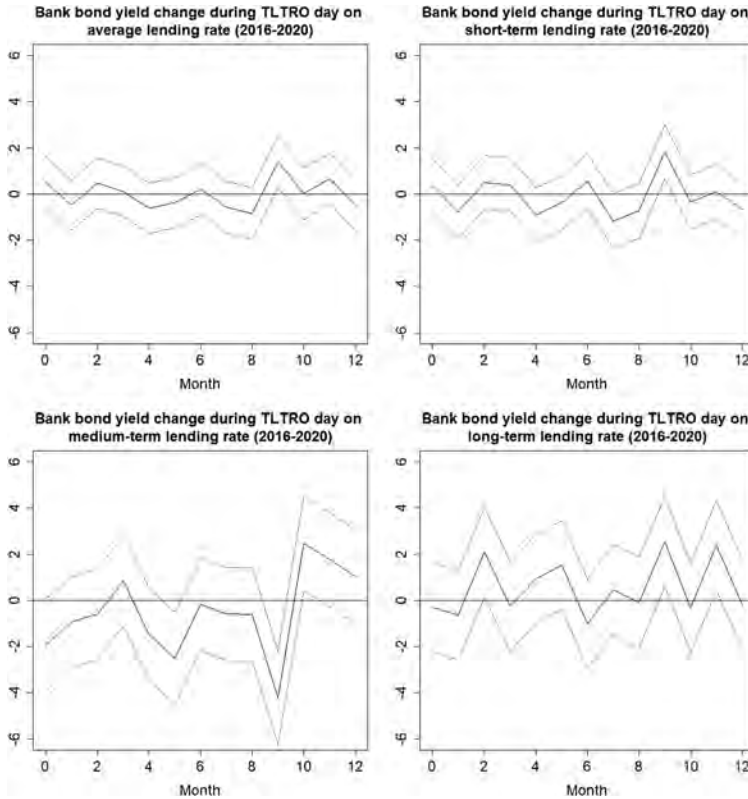


Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

an actual measure of each bank’s amount of funds lent from these longer-term operations, so we can test this directly. We thus augment the model by an interaction term with our (T)LTRO measure and the lagged share of central bank funding on bank *i*’s balance sheet. Figure 13 shows the results for the 2014–20 period.

We see that the interaction is positive and significant, especially when it comes to loans of long maturities. The result is intuitive, as

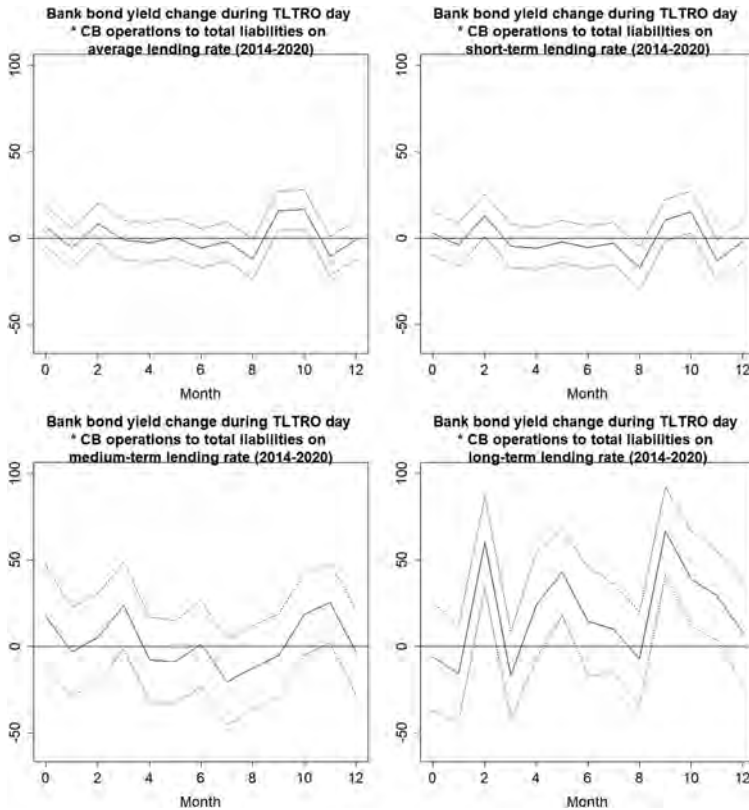
Figure 12. Effect of (T)LTRO Shock for the Average Corporate Lending Rate and Three Different Maturities, Low-for-Long Period (2016–20)



Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

the TLTROs had a built-in incentive for lending in a certain time window. The banks were rewarded with a lower interest rate by the ECB if lending was increased during this relatively long monitoring period. It thus made less sense to lend short-term, as the ECB only cared about the average lending change at the end of the period. For example, in the case of TLTRO II, launched in 2016, the potential reward to banks for lending depended on their total growth in

Figure 13. Effect of (T)LTRO Shock on Banks with High Share of TLTRO Loans for the Average Corporate Lending Rate and Three Different Maturities, Period of Negative Rates (January 2014–December 2020)



Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

lending (loans to firms and households, excluding lending for house purchases) from February 2016 to January 2018. Therefore, there were fewer incentives for lending short-term (at least immediately after the announcement). We conclude that TLTROs contributed to mitigating the pass-through of monetary policy easing to bank lending rates below zero, and this effect was stronger for banks taking on more funding directly from these operations.

5.3 *Asset Purchase Programs*

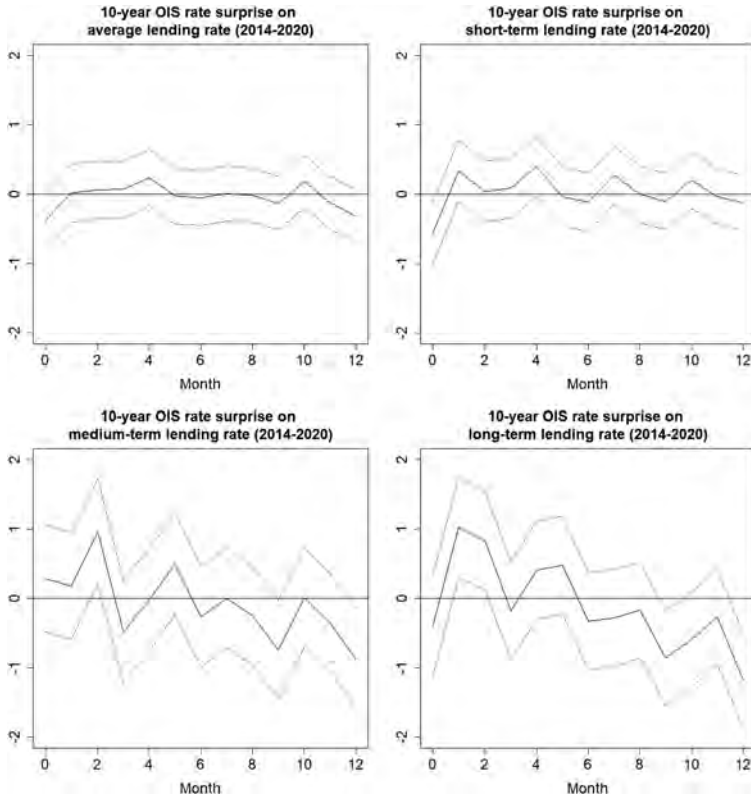
Finally, we focus on the ECB large-scale asset purchase programs, or QE, proxied by the 10-year OIS rate shock. We include in Equation (1) the one-week OIS rate surprise as a contemporaneous control variable. This control is added because, as conventional monetary policy may affect long-term rates, asset purchases do not affect short-term rates contemporaneously. For QE, we only have observations starting in 2014, when the ECB launched the APP package and other large-scale asset purchase programs. The estimations for the full period of 2010–20 are thus omitted in this subsection.

Figure 14 presents the results for the period of negative rates. The results suggest that QE had no effect on average corporate lending rates; the coefficients are statistically insignificant. However, there is evidence that QE had some positive effect on lending rates of medium- and long-term loans at shorter horizons. That is, as QE induced an unexpected decrease at the longer end of the yield curve, banks lowered their lending rates for loans of longer maturities. This could stem from the fact that they are more likely to be closer substitutes for long-term bonds. Focusing solely on the low-for-long period (Figure 15), we find no clear evidence of a favorable effect on lending rates.

Bank heterogeneity can again affect the impact of QE. With asset purchases having a large impact on security prices, the market-to-market value of bank security holdings increases, raising bank net worth (Brunnermeier and Sannikov 2014). Assuming that commercial banks target somewhat constant leverage ratios, this induces banks to expand their lending (Adrian and Shin 2010), which here would be related to lower lending rates. To account for bank heterogeneity, we use bank-level data on bond holdings. Figure 16 shows the estimated coefficients for the interaction term between the 10-year rate surprise and the lagged share of bonds in a bank's total assets. Our results, however, provide no evidence that QE policies would have any stronger effect on lending rates of banks directly more exposed to bond purchases.

The ABSPP and CBPP3 asset purchase programs, which focus on asset-backed securities and covered bonds, are linked to loans to the private sector granted by banks—not to the bonds they are holding. Although their share is small compared to the PSPP, focusing

Figure 14. Effect of QE Shock for the Average Corporate Lending Rate and Three Different Maturities, Period of Negative Rates (January 2014–December 2020)



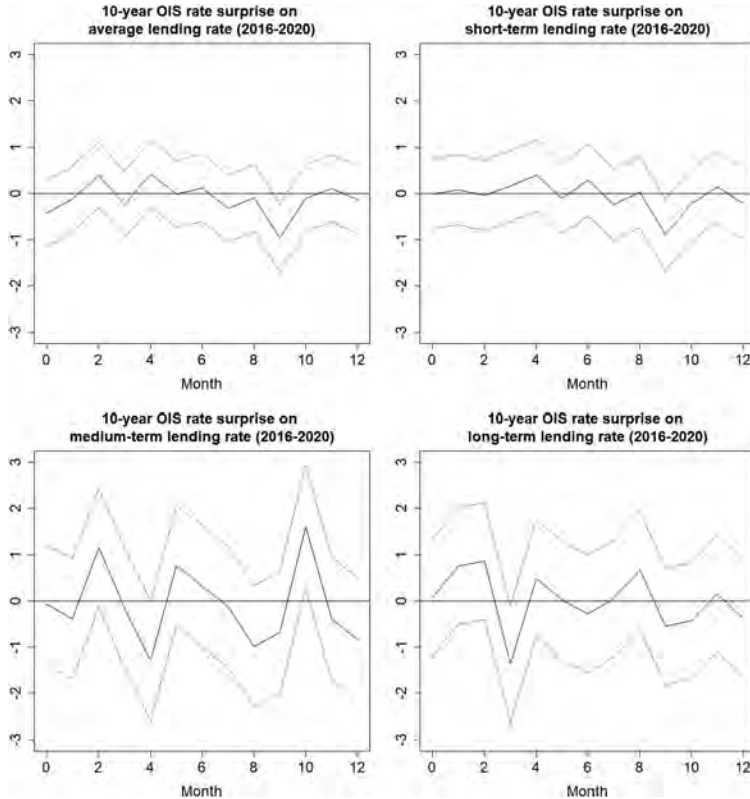
Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

on public-sector bonds, this could partly explain why banks specifically holding more bonds are not differently affected by QE shocks in our estimations.

6. Conclusion

In this paper, we study how various monetary policy measures adopted by the ECB transmit to corporate lending rates in a

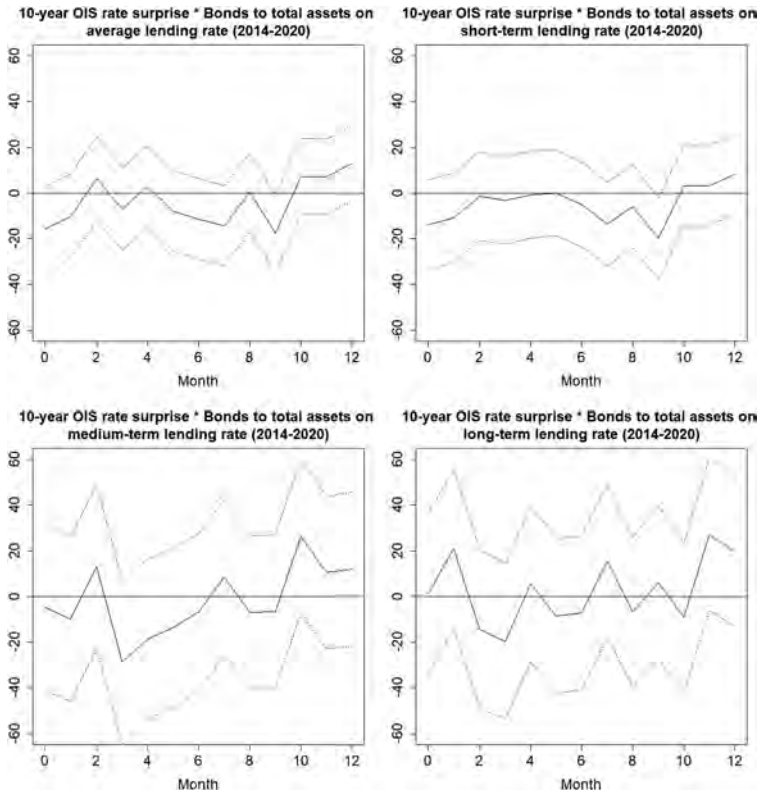
Figure 15. Effect of QE Shock for the Average Corporate Lending Rate and Three Different Maturities, Low-for-Long Period (January 2016–December 2020)



Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

negative interest rate environment. Utilizing a detailed bank-level data set covering 137 individual banks from 13 euro-area countries at a monthly frequency for the period from January 2010 to December 2020, we examine how the pass-through changed when policy rates turned negative in June 2014 and further after early 2016, when negative rates came to be regarded as persistent due to the ECB's revised forward guidance and additional rate cuts.

Figure 16. Effect of QE Shock on Banks with High Share of Bond Holdings for the Average Corporate Lending Rate and Three Different Maturities, Period of Negative Rates (January 2014–December 2020)



Note: Maturities: less than one year, one to five years, and over five years. Time horizon: 12 months. 90 percent confidence intervals à la Driscoll and Kraay (1998) are reported.

We find several noteworthy results. First, the transmission of short-term policy rates is clearly hampered below zero. Although we already see some signs of the reversal rate during the 2014–20 period, the evidence becomes much stronger as negative rates become more persistent during the low-for-long period. Following further policy rate cuts, banks started to raise their lending rates. Second, bank


heterogeneity plays a role. Stronger evidence for the reversal rate is obtained for banks more exposed to negative policy rates: those that did not lower their own retail deposit rates below zero, and those holding more negative interest-bearing central bank deposits on their balance sheets. Third, loan maturities matter. The inverse peak effect occurs at different horizons for different maturities, which may explain the weaker results when looking at the average lending rates. Moreover, the statistical significance and the absolute size of the inverse effect increase with the length of the loan. Fourth, even if the transmission of the short-term policy rate to bank lending rates is hampered below zero, unconventional monetary policy measures—especially TLTROs—help mitigate the pass-through by lowering bank lending rates, especially for loans of longer maturities.

Our analysis helps with understanding why empirical studies so far have struggled to find common ground for the existence of the reversal rate. The results confirm what previous literature suggests: that its existence depends greatly on bank heterogeneity, such that the reversal rate is more pronounced for banks that are most exposed to negative rates (those that are reluctant to lower their own retail deposit rates below zero or those with large amounts of negative interest-bearing central bank deposits). We also show that different loan maturities play an important role, as the inverse effect is stronger for loans of longer maturities. Furthermore, by employing a long enough data sample with negative rates, we confirm that negative rates must first become persistent before we can seek to uncover evidence of the reversal rate.

Looking to the policy implications of these findings, we provide evidence that monetary policy rate cuts cease to have the desired effect on private-sector lending costs when negative rates become persistent. Without the mitigating effects of additional policies specifically aimed at lowering bank funding costs, further rate cuts below zero run the risk of becoming contractionary. Moreover, although reducing volatility and uncertainty about future interest rate developments, central bank signaling of a low-for-long environment could bring the reversal rate forward.

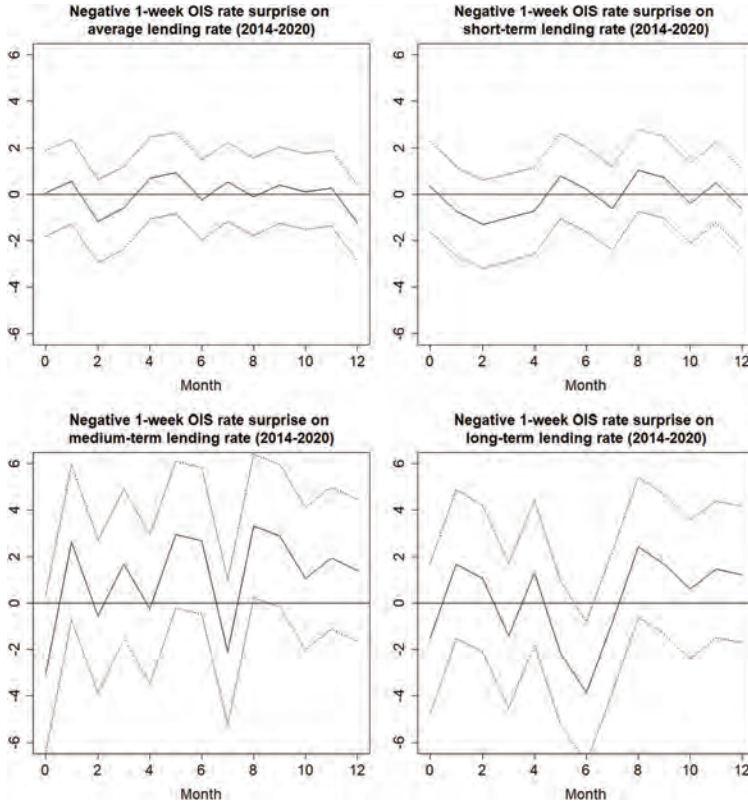
Appendix

Figure A.1. Representativeness of Our Sample



Note: Euro-area totals: ECB MIR and BSI databases.

Figure A.2. Results Regarding the Short-Term Rate Surprise for the Negative Rates Sample (2014–20) When Only Negative Rate Surprises Are Considered



Note: The results show that rate cut surprises in the below-zero environment increase lending rates. Note that the figure reports regression coefficients as such (not multiplied by -1).

Table A.1. Rolling Window Results

Subsample/ Horizon	2010–14	2011–15	2012–16	2013–17	2014–18	2015–19	2016–20
0	0.01	0.07	0.63	1.03	0.13	0.11	0.86
1	0.15	0.44	1.60***	1.68***	1.28***	1.37*	0.56
2	0.01	0.34	0.54	-0.35	-1.58***	-1.04**	-0.23
3	0.65**	0.67**	1.25**	0.46	-0.09	-0.89	-0.55
4	0.28	0.43	0.53	0.25	0.03	-0.15	0.12
5	-0.34	-0.35	0.5	0.87	1.03**	-1.67**	-0.13
6	0.34	0.13	0.05	-0.28	-0.56	-0.66	0.23
7	0.4	-0.02	-0.36	-0.4	0.25	-1.00*	-3.43***
8	-0.15	-0.56	-0.78	-0.96**	-0.66**	-0.31	-0.69
9	1.10**	0.22	-0.95***	-0.87*	0.53	0.31	-2.85
10	0.42	-0.34	-1.08***	-0.73	-0.22	0.79	0.96
11	0.7	0.04	0.33	0.21	0.47	-0.06	-3.07**
12	0.13	-0.52	-1.28***	-0.62	-0.81	1.25	0.21

Note: Monetary policy is measured using the one-week OIS surprise. The endogenous variable is the average lending rate. The models include our standard set of control variables that are used in all the local projection analyses: GDP growth, inflation, and unemployment forecasts from the Survey of Professional Forecasters (SPF) to control for expectations about the macroeconomic environment, as well as growth in industrial production, the unemployment rate, and the Eurostoxx stock return to control for the current environment. Bank-specific controls include liquidity, capitalization, and bank size. We control for credit demand using dummy variables based on the bank lending survey. In addition, we control for the stance of monetary policy prior to the shock using the EONIA and Krippner shadow rates. Three lags of each control variable are included. The standard errors are calculated based on a nonparametric robust covariance matrix estimator à la Driscoll and Kraay (1998). *, **, *** denote an estimate significantly different from zero at the 10 percent, 5 percent, and 1 percent level, respectively.

Table A.2. Statistically Significant Peak Effects on Average Lending Rates of a One-Unit (100 Basis Point) Surprise in the Shock Proxy of Interest

Subsample/Policy Tool	2010–20	2014–20	2016–20
Conventional MP	1.07 (h=9)	-1.07 (h=2)	-3.43 (h=7)
QE		-0.38 (h=0)	-0.96 (h=9)
TLTRO	1.82 (h=6)	0.72 (h=9)	1.41 (h=9)
Note: Horizon (h) is reported in parentheses.			

Table A.2 summarizes our main results by reporting the statistically significant peak effects on average lending rates of a one-unit (100 basis point) surprise in the shock proxy of interest.

Figure A.3.

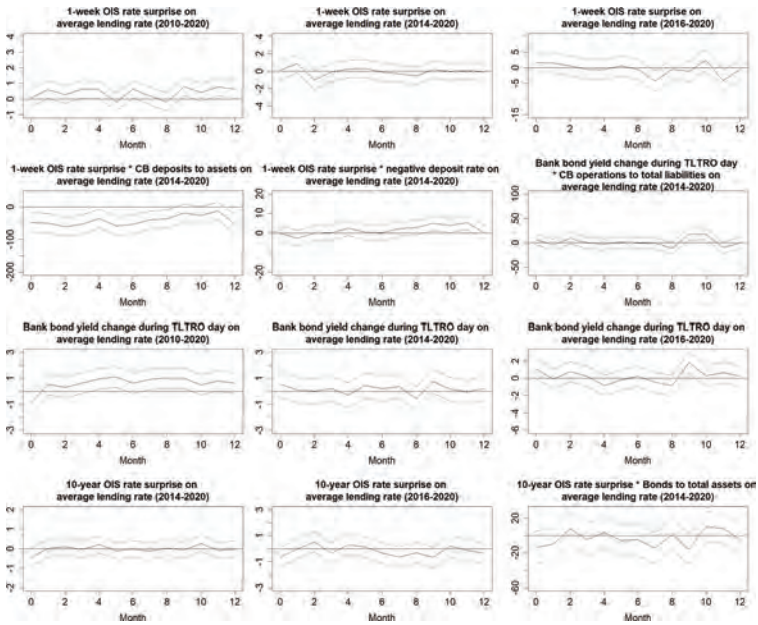


Figure A.3 presents our results regarding average lending rates after controlling for the leads ($t+3$) of SPF GDP growth expectation and loan demand dummy variables.

Figure A.4.

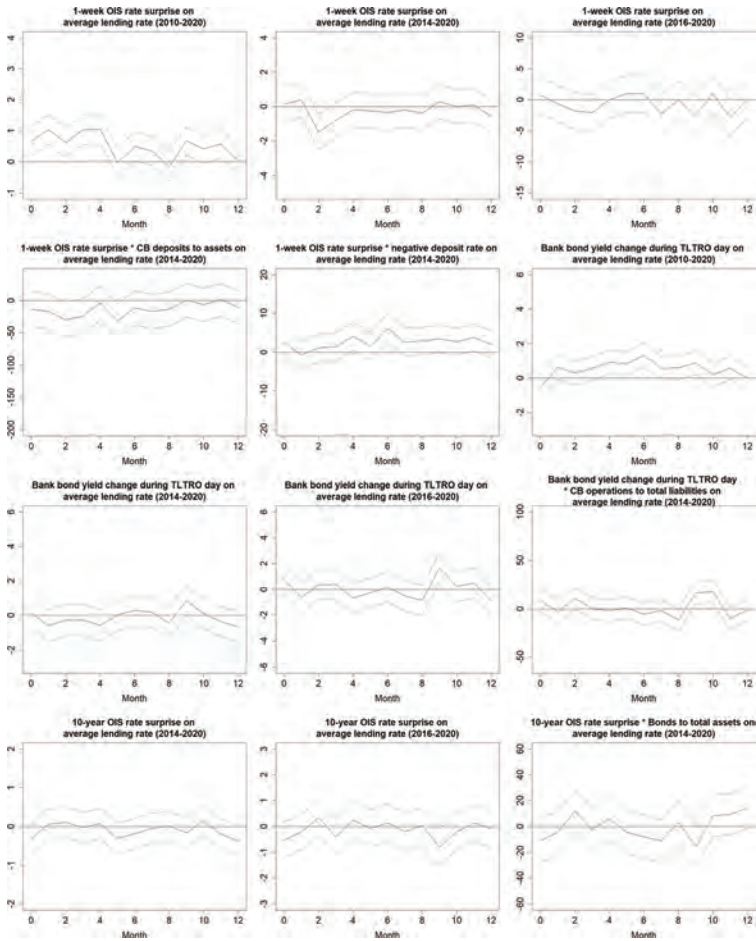


Figure A.4 presents the results for the average lending rates after replacing country-specific macro variables with the country and country-year fixed effects.


Figure A.5.

Figure A.5 presents the results regarding conventional monetary policy for the cumulative change of (log) loan amounts from $t-1$ to $t+h$.

References

- Abadi, J., M. K. Brunnermeier, and Y. Koby. 2023. “The Reversal Interest Rate.” *American Economic Review* 113 (8): 2084–120.
- Adrian, T., and H. S. Shin. 2010. “Liquidity and Leverage.” *Journal of Financial Intermediation* 19 (3): 418–37.
- Albertazzi, U., A. Nobili, and F. M. Signoretti. 2021. “The Bank Lending Channel of Conventional and Unconventional Monetary Policy.” *Journal of Money, Credit and Banking* 53 (2–3): 261–99.
- Altavilla, C., F. Barbiero, M. Boucinha, and L. Burlon. 2023. “The Great Lockdown: Pandemic Response Policies and Bank Lending Conditions.” *European Economic Review* 156 (July): 104478.
- Altavilla, C., M. Boucinha, and P. Bouscasse. 2022. “Supply or Demand: What Drives Fluctuations in the Bank Loan Market?” ECB Working Paper No. 2646.
- Altavilla, C., L. Brugnolini, R. Gürkaynak, R. Motto, and G. Ragusa. 2019. “Measuring Euro Area Monetary Policy.” *Journal of Monetary Economics* 108: 162–79.
- Altavilla, C., L. Burlon, M. Giannetti, and S. Holton. 2022. “Is There a Zero Lower Bound? The Effects of Negative Policy Rates on Banks and Firms.” *Journal of Financial Economics* 144 (3): 885–907.

- Ampudia, M., M. Ehrmann, and G. Strasser. 2023. "The Effect of Monetary Policy on Inflation Heterogeneity along the Income Distribution." ECB Working Paper No. 2858.
- Andrade, P., C. Cahn, H. Fraise, and J. S. Mésonnier. 2019. "Can the Provision of Long-Term Liquidity Help to Avoid a Credit Crunch? Evidence from the Eurosystem's LTRO." *Journal of the European Economic Association* 17 (4): 1070–106.
- Balloch, C., Y. Koby, and M. Ulate. 2022. "Making Sense of Negative Nominal Interest Rates." Federal Reserve Bank of San Francisco Working Paper No. 2022-12. <https://doi.org/10.24148/wp2022-12>.
- Basten, C., and M. Mariathasan. 2023. "Interest Rate Pass-Through and Bank Risk-Taking under Negative-Rate Policies with Tiered Remuneration of Central Bank Reserves." *Journal of Financial Stability* 68: 101160.
- Benetton, M., and D. Fantino. 2021. "Targeted Monetary Policy and Bank Lending Behavior." *Journal of Financial Economics* 142 (1): 404–29.
- Bluedorn, J., C. Bowdler, and C. Koch. 2017. "Heterogeneous Bank Lending Responses to Monetary Policy: New Evidence from a Real-Time Identification." *International Journal of Central Banking* 13 (1): 95–149.
- Boeckx, J., M. de Sola Perea, and G. Peersman. 2020. "The Transmission Mechanism of Credit Support Policies in the Euro Area." *European Economic Review* 124 (May): 103403.
- Boeckx, J., M. Dossche, and G. Peersman. 2017. "Effectiveness and Transmission of the ECB's Balance Sheet Policies." *International Journal of Central Banking* 13 (1): 297–333.
- Borio, C., and L. Gambacorta. 2017. "Monetary Policy and Bank Lending in a Low Interest Rate Environment: Diminishing Effectiveness?" *Journal of Macroeconomics* 54: 217–31.
- Bottero, M., C. Minoiu, J.-L. Peydró, A. Polo, A. Presbitero, and E. Sette. 2022. "Expansionary yet Different: Credit Supply and Real Effects of Negative Interest Rate Policy." *Journal of Financial Economics* 146: 754–78.
- Brunnermeier, M. K., and Y. Sannikov. 2014. "A Macroeconomic Model with a Financial Sector." *American Economic Review* 104 (2): 379–421.

- Claessens, S., N. Coleman, and N. Donnelly. 2018. “Low-for-Long’ Interest Rates and Banks’ Interest Margins and Profitability: Cross-Country Evidence.” *Journal of Financial Intermediation* 35: 1–16.
- Cox, D. R. 1961. “Prediction by Exponentially Weighted Moving Averages and Related Methods.” *Journal of the Royal Statistical Society Series B* 23 (2): 414–22.
- Demiralp, S., J. Eisenschmidt, and T. Vlassopoulos. 2021. “Negative Interest Rates, Excess Liquidity and Retail Deposits: Banks’ Reaction to Unconventional Monetary Policy in the Euro Area.” *European Economic Review* 136 (July): 103745.
- Driscoll, J. C., and A. C. Kraay. 1998. “Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data.” *Review of Economics and Statistics* 80 (4): 549–60.
- Eggertsson, G. B., R. E. Juelsrud, L. H. Summers, and E. G. Wold. 2023. “Negative Nominal Interest Rates and the Bank Lending Channel.” *The Review of Economic Studies* 91 (4): 2201–75. <https://doi.org/10.1093/restud/rdad085>.
- Erikson, H., and D. Vestin. 2019. “Pass-Through at Mildly Negative Policy Rates: The Swedish Case.” Staff Memorandum, Sveriges Riksbank, January.
- European Central Bank. 2016. “Introductory Statement to the Press Conference (with Q&A).” Mario Draghi, President of the ECB and Vitor Constâncio, Vice-President of the ECB, Frankfurt am Main, January 21. https://www.ecb.europa.eu/press/press_conference/monetary-policy-statement/2016/html/is160121.en.html.
- Gertler, M., and N. Kiyotaki. 2010. “Financial Intermediation and Credit Policy in Business Cycle Analysis.” *Handbook of Monetary Economics* 3: 547–99.
- Heider, F., F. Saidi, and G. Schepens. 2021. “Banks and Negative Interest Rates.” *Annual Review of Financial Economics* 13: 201–18.
- Holston, K., T. Laubach, and J. C. Williams. 2017. “Measuring the Natural Rate of Interest: International Trends and Determinants.” *Journal of International Economics* 108: S59–S75.
- Jarociński, M., and P. Karadi. 2020. “Deconstructing Monetary Policy Surprises—The Role of Information Shocks.” *American Economic Journal: Macroeconomics* 12 (2): 1–43.

- Jordà, Ò. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–82.
- Kho, S. 2025. "Deposit Market Concentration and Monetary Transmission: Evidence from the Euro Area." *European Economic Review* 173 (April). <https://doi.org/10.1016/j.euroecorev.2024.104933>.
- Kiley, M. T., and J. M. Roberts. 2017. "Monetary Policy in a Low Interest Rate World." *Brookings Papers on Economic Activity* 2017 (1): 317–96.
- Krippner, L. 2015. *Zero Lower Bound Term Structure Modeling: A Practitioner's Guide*. Springer.
- Krishnamurthy, A., and A. Vissing-Jorgensen. 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." *Brookings Papers on Economic Activity* 43 (2): 215–87.
- Kwan, S., M. Ulate, and V. Voutilainen. 2025. "The Transmission of Negative Nominal Interest Rates in Finland." *Journal of the European Economic Association* 23 (5): 1809–37. <https://doi.org/10.1093/jeea/jvaf005>.
- Laine, O.-M. 2021. "The Effect of Targeted Monetary Policy on Bank Lending." *Journal of Banking and Financial Economics* 1 (15): 25–43.
- Laine, O.-M., and M. Pihlajamaa. 2024. "Pushing and Pulling on a String? Inflationary Effects of Expansionary and Contractionary Monetary Policies When Rates Are Negative." *Economic Modelling* 131: 106620.
- Molyneux, P., A. Reghezza, J. Thornton, and R. Xie. 2019. "Did Negative Interest Rates Improve Bank Lending?" *Journal of Financial Services Research* 57 (1): 51–68.
- Plagborg-Møller, M., and C. K. Wolf. 2021. "Local Projections and VARs Estimate the Same Impulse Responses." *Econometrica* 89 (2): 955–80.
- Rostagno, M., C. Altavilla, G. Carboni, W. Lemke, R. Motto, A. S. Guilhem, and J. Yiangou. 2021. *Monetary Policy in Times of Crisis: A Tale of Two Decades of the European Central Bank*. Oxford University Press.
- Shimko, D. 1993. "Bounds on Probability." *Risk* 6 (4): 33–37.

- Ulate, M. 2021. "Going Negative at the Zero Lower Bound: The Effects of Negative Nominal Interest Rates." *American Economic Review* 111 (1): 1–40.
- Wang, Y., T. Whited, Y. Wu, and K. Xiao. 2022. "Bank Market Power and Monetary Policy Transmission: Evidence from a Structural Estimation." *Journal of Finance* 77 (4): 2093–141.

Modernizing Central Banks' Large Semi-Structural Projection Models: The Role of Model-Consistent Expectations*

Stéphane Adjemian,^a Nikola Bokan,^b
Matthieu Darracq Pariès,^b Georg Müller,^b Srećko Zimic^b

^aUniversité du Mans

^bEuropean Central Bank

We enrich the ECB's modeling toolkit by incorporating model-consistent expectations into the workhorse projection model, ECB-BASE. This enables analysis of announced policy shifts and anticipation effects. We nest VAR expectations, model-consistent expectations and several combinations of the two, termed as hybrid settings. In addition, the Dynare computational environment has been advanced in several dimensions to enable these new modeling capabilities. We then examine monetary policy implications of varying expectation formation. Simulations indicate that forward-looking expectations result in smoother output responses but front-loaded price reactions to policy shocks. By applying the model to assess recent supply shocks that have led to a surge in inflation, we find that forward-looking expectations from the private sector imply increased macroeconomic volatility, necessitating an early policy response.

JEL Codes: C3, C5, E1, E2, E5.

1. Introduction

This paper contributes by modernizing the semi-structural projection model toolkit of the European Central Bank (ECB). We study the transmission of monetary policy in the ECB-(RE)BASE model

* Author email addresses: stephane.adjemian@univ-lemans.fr (S. Adjemian), nikola.bokan@ecb.europa.eu (N. Bokan), matthieu.darracq-paries@ecb.europa.eu (M. Darracq Pariès), georg.muller@ecb.europa.eu (G. Müller, corresponding author), srecko.zimic@ecb.europa.eu (S. Zimic).

across different expectation formation alternatives. ECB-(RE)BASE is a model-consistent expectations (MCE), or rational expectations (RE), version of the ECB-BASE model, which up until the time of writing operated solely under expectations formed via vector autoregression (VAR). The underlying ECB-BASE model developed at the ECB was introduced by Angelini et al. (2019) as the primary semi-structural macroeconomic model for the euro area. Today, it is instrumental in informing ECB monetary policy preparation by contributing to economic projections, risk and scenario analysis, and quantitative policy evaluations. The main advantage of such models lies in the balance between the theoretically founded structure (and its imposed restrictions) and the empirical validity of captured gradual transmissions. Importantly, expectations also play a critical role in determining macroeconomic dynamics in this class of models. As a functional policy analysis tool, the model features a comprehensive description of the euro-area economy, incorporating a detailed demand breakdown, price and wage developments, forward-looking asset pricing, financing conditions, and future income projections, as well as different roles for policy interventions.

The impetus to explore alternative expectation settings arises from recent supply shocks that have increased energy, production, and consumer prices. These shocks have heightened macroeconomic volatility, underscoring the need to understand the role of expectations as long-term anchors. Induced by such developments, the monetary authorities needed to assess the implications of different degrees of forward-lookingness on macroeconomic volatility. After a prolonged period of low interest rates, the central bank's reaction to the new environment resulted in many instances in a large hiking cycle. In this respect, our work in this paper focuses on understanding the consequences of unexpected and anticipated monetary policy shocks.¹

At the same time, central bank forecasting models are under increasing scrutiny (see Bernanke 2024). Dynamic stochastic general equilibrium (DSGE) models typically emphasize forward-looking behavior under the rational expectations paradigm. In contrast,

¹The accompanying working paper by Adjemian, Bokan et al. (2024) systematically documents the implications of different expectation settings for other main shocks.

the majority of semi-structural projection models rely predominantly on established long-run relationships consistent with data. For pure forecasting purposes, it is often adequate to consider expectations based on historical patterns. However, for analytical purposes, exploring alternative expectation settings that account for policy shifts, evaluate announced changes, and anticipate news would be highly beneficial. We contribute to the ECB's analytical capabilities by integrating these alternative expectation settings into the main projection tool. The implementation also allows for hybrid settings, combining different expectations across various sectors or blocks of the model simultaneously.

Changing expectations from backward-dependent to forward rational expectations affects monetary policy transmission. We find that using the MCE setting leads to a dampened cyclical response because of stronger income smoothing and more gradual transition dynamics. These smoothing forces dominate any implications of the financial sector being forward-looking. We also find that there are front-loading features on prices, especially on HICP (Harmonised Index of Consumer Prices) inflation through the exchange rate effect on imported energy, but less on the GDP deflator. We show that policy announcements lead to partial anticipation of economic dynamics if credibly understood by the private sector. Additionally, we show that recent supply shocks imply more macroeconomic volatility under MCE compared to VAR expectations, highlighting the need for central banks to understand the degree of anticipation of supply pressures to achieve their stabilization objectives.

In its primary role as a forecasting and policy analysis tool, the ECB-BASE model operates under VAR expectations and is estimated accordingly. In this setup, current and past states potentially affect expectations. The VAR provides a reduced-form view of the full economic environment described by the complete macro model and can, loosely speaking, be considered as a "limited or condensed information" expectation formation version. Other forms of restricted information sets used in the forecasting literature are adaptive learning approaches (Slobodyan and Wouters 2012; Warne 2023).

A key challenge in swapping expectations in this class of models is the adjustment costs. These are central for the model to generate empirically plausible transition dynamics. They are modeled via the

polynomial adjustment costs (PAC) approach, which is an error-correction-type specification with an explicit role of forward-looking expectations. The Dynare computational environment has been extended to support the reformulation of a VAR-based setup to the MCE setup.

This paper also relates to the literature on different expectation formation models, in particular in the context of semi-structural models as used in central banks. The main reference point is the work around the FRB/US model (Brayton and Tinsley 1996; Brayton et al. 1997), which served as inspiration in the development of the underlying ECB-BASE model. In their analytical work, Federal Reserve staff have also occasionally compared the impulse responses the model produces under different expectation settings. Brayton, Laubach, and Reifschneider (2014) show that VAR-based expectations can imply high inertia that differs from the inertia implied by the Taylor rule. In this case, the model-consistent expectations imply a softer price response to a monetary policy rate shock than under VAR expectations.

Numerous FRB/US simulations with an emphasis on the role of expectations have been used for monetary policy analysis (e.g., Reifschneider and Williams 2000; Chung et al. 2012, among others). Kiley and Roberts (2017) evaluate the frequency and potential costs of Lower Bound incidences, showing that clear commitment strategies by the central bank to maintain lower interest rates until it achieves the inflation objective are highly effective under MCE in the FRB/US model. Bernanke, Kiley, and Roberts (2019) find that MCE can imply lower output volatility than hybrid expectations when there is a risk of hitting the Lower Bound, as private-sector expectations regarding the inflation anchor help alleviate the constraint. Tetlow (2022) uses FRB/US to assess the costs of disinflation, showing that a persistent reduction in inflation can be achieved under MCE via a credible announcement of shifting the inflation anchor without sharply decreasing demand. Conversely, VAR expectations require a more aggressive policy adjustment, causing a more significant decline in demand to achieve the same outcome.

The remainder of our paper is structured as follows. Section 2 lays out the strategy to solve a model under rational expectations. It describes the role of expectations in ECB-(RE)BASE while broadly outlining its main features and shows the details necessary to convert

BASE into a forward-looking expectation model. Section 3 examines impulse responses to monetary policy shocks, comparing different expectation modalities, and discusses some key contributing factors. Section 4 examines the inflation surge across different expectation settings. Finally, Section 5 concludes.

2. ECB-(RE)BASE

This section discusses the modifications to the ECB-BASE model to convert it into a model-consistent expectations variant.

Before delving into the model details, it is useful to generally recall the treatment of expectations in solving dynamic macro models. Consider the following type of model:

$$\mathbb{E}_t[f(y_{t+1}, y_t, y_{t-1}, \epsilon_t)] = 0, \quad (1)$$

where y_t is a vector of endogenous variables, ϵ_t is a vector of structural shocks with $\mathbb{E}_t[\epsilon_{t+1}] = 0$ and $\mathbb{E}_t[\epsilon_{t+1}\epsilon'_{t+1}] = \Sigma$. Under backward expectations, the terms $\mathbb{E}_t[y_{t+1}]$ are replaced by² VAR forecasts: $y_{t+1} = h'x_t$, where x_t is a set of explanatory variables that may include y_t and lags thereof. The solution then amounts to solve for y_t , in each period given $\underline{y_{t-1}}$ and ϵ_t . Formally, the following nonlinear problem arises:

$$f(h'x_t, y_t, y_{t-1}, \epsilon_t) = 0. \quad (2)$$

In this case, the paths for the endogenous variables are the unknown object of interest that we want to solve for. Contrary to that strategy, under rational expectations, we solve Equation (1) for an invariant mapping between the endogenous variables and the state variables. Let this mapping be $y_t = g(y_{t-1}, \epsilon_t)$. We can then reformulate Equation (1) as

$$\mathbb{E}_t[f(g(y_{t-1}, \epsilon_t), \epsilon_{t+1}), g(y_{t-1}, \epsilon_t), y_{t-1}, \epsilon_t)] = 0, \quad (3)$$

²Assuming the model exhibits linearity with respect to the forward terms, the conditional expectation can be accurately transferred under the function f . In contrast, for nonlinear models, this operation results in an approximation due to Jensen's inequality. Semi-structural models are typically linear with respect to the forward variables.

which is typically solved by perturbation techniques around the deterministic steady state. If we consider a first-order approximation, the solution is a linear reduced-form model: $y_t = Ay_{t-1} + B\epsilon_t$.

Let us now recall the basics of the ECB-BASE model. Figure 1 gives a bird's-eye view of the model structure. The rich structure of the model enables us to incorporate various mechanisms behind economic forces and policy channels. More specifically, the wage-price block is based on a simple auxiliary New Keynesian model used to estimate the model's main price and wage Phillips curves (see Appendix B.4 in Angelini et al. 2019).

The model's domestic demand structure features private consumption, business and housing investment, and government expenditures. Global forces, captured in the foreign block, affect exports and imports. Financing conditions, as well as asset prices, influence borrowing and transmit into the consumption and investment decisions.

In the wage-price block, reduced-form pricing curves assign an important role to forward-looking expectations.³ Important to notice is that this auxiliary model is estimated under rational expectations as a stand-alone system. It is therefore straightforward to revert to this specification when transitioning from a VAR to an MCE setup.⁴ Given the prominence of the financial block in the model, in (RE)BASE, we respecify the model's asset pricing problem and align it to the formulations standard to the DSGE literature.⁵ To do so, it is convenient to change the logic of how bond prices and financing rates are constructed. Furthermore, we also treat the exchange rate mechanism as a financial feature. Postulating the exchange rate as

³The reduced-form goods price Phillips curve is of the form $\pi_t = \{(1 - \delta_\pi)(1 - \beta_\pi)(1 - \varphi_\pi)\}^{-1}\{(1 - \delta_\pi)(1 - \beta_\pi)(1 - \varphi_\pi)\bar{\pi}_t^e + \beta_\pi E_t \pi_{t+1} + \delta_\pi \pi_{t-1} + \beta_y^\pi mc_t + \varphi_\pi \pi_t^{oil}\} + \epsilon_t^\pi$. π_t is the core deflator inflation, mc_t are marginal costs, and π_t^{oil} is imported energy price inflation. The medium-term inflation expectations, $\bar{\pi}_t^e$, are imperfectly anchored and are updated based on past inflation realizations. They follow this rule throughout this paper.

⁴Given that (RE)BASE and BASE share an identical specification of the wage-setting structure, we do not restate it here.

⁵House prices are modeled as in ECB-BASE. They are not per se modeled as financial asset valuations like the bond prices are, but rather are determined according to simple empirical relationships based on contemporaneous information and on the medium-term inflation anchor. However, the transition dynamics of house prices to the target value follow the PAC framework.

Figure 1. Overview of the ECB-BASE Model and the Center-Stage of Expectations



Note: Each small hexagon represents a specific model “block.” The blocks gravitate toward specific long-run anchors, which discipline the model through the expectation terms. In the case of VAR expectations, a core VAR serves to pin down the main relationships between inflation, the output gap, and the interest rate, which is then extended to account for other explanatory forces differing across model blocks, while at the same time ensuring a certain degree of system-wide consistency. The large light blue shaded hexagon indicates where these expectations enter. Furthermore, the main adjustment dynamics governed by polynomial adjustment costs (PAC) are also subjected to the expectations. The framed light blue hexagon touches the blocks where PAC play a role. To be precise, they govern the transition dynamics of private consumption, private business and residential investment, house prices, property income, and employment.

a forward-looking variable implies a significant departure from the empirical setup used in the original model.

An interesting specificity of the model is the prominence of multiple income sources (labor, transfer, and property income) as main

determinants of consumption. While the underlying theory behind the relationships entering the VAR expectations and MCE version of the model is the same, the former requires significant algebraic manipulation of the equations in order to arrive at a trend-gap representation. On the other hand, the MCE version of the model eliminates the need for this trend-gap decomposition. Therefore, in order to facilitate the exposition of the changes in the model structure versus the original BASE, we do present the aforementioned changes in more detail in the following section.

We should also reiterate the significant role played by the expectations in the adjustment process of the model's main macroeconomic variables. The adjustment to the long-run trend values is a sluggish transition process arising as a consequence of the various adjustment costs featured in the model and governed, at least in part, by the expectations. Figure 1 illustrates the various blocks featuring expectations and their main formation mechanism. Two light blue shaded hexagonal areas overlapping each other connect the model blocks where either core VAR expectations and/or PAC-based formulations enter the model.

In the following, we describe the modifications in the model necessary to swap it into its MCE form. The order of exposition pre-emptly the order in which we want to switch on certain features sequentially in the simulation exercises presented later.

2.1 Modified Financial Blocks

2.1.1 Bond Prices and the Long-Term Interest Rate

The following assets are modeled with forward-looking expectation formation: 10-year (long-term) bonds, corporate bonds, and equities/stocks. The short-term nominal bond price is not explicitly modeled. The policymaker steers the value of short-term bonds by setting a nominal interest rate according to the Taylor rule.

The long-term interest rate in ECB-(RE)BASE is a 10-year benchmark government bond yield. Computationally, it is convenient to use infinitely-lived consol-style bond pricing equations, which have a tractable recursive formulation. We follow the approach of Woodford (2001) such that the long-term bond is modeled as a

coupon-paying consol.⁶ Coupon payments are expected to occur in each period, starting from $t+1$, and the duration of the bond is chosen to reflect a 10-year horizon of payments.⁷

To price the long-term bond, investors discount expected future coupon payments and add a term premium, $\bar{t}p$, to the one-period-ahead pricing kernel. The bond price, Q_t^{10Y} , is given by the following recursive forward form:

$$Q_t^{10Y} = \frac{c}{(1 + i_t)} + \rho \mathbb{E}_t \left[\frac{Q_{t+1}^{10Y}}{(1 + i_t)(1 + \bar{t}p)} \right], \tag{4}$$

where i_t is the short-term nominal interest rate, c is the coupon payment, calibrated to normalize a steady-state bond price, $\overline{Q^{10Y}}$, to 1, and ρ is the decaying factor of coupon payments, calibrated to match a Macaulay duration of 40 quarters. $\bar{t}p$ is the term premium associated with holding a bond for a 10-year period.⁸

The long-term interest rate, i_t^{10Y} , is then defined as the compounded yield to maturity on the consol,

$$i_t^{10Y} = c/Q_t^{10Y} + \rho - 1. \tag{5}$$

Note that the rate of return on bond holding r_t^{10Y} is given by

$$r_t^{10Y} = \frac{c + \rho Q_t^{10Y}}{Q_{t-1}^{10Y}} - 1. \tag{6}$$

A similar modeling strategy is adopted for the value of corporate bonds and its rate of return, r_t^{CB} . It follows a specification as

⁶In the original ECB-BASE model, in contrast, bond values are not explicitly modeled. Instead, the risk-free interest rate according to the expectation hypothesis is formed based on a VAR forecast. A term premium is then added to reflect the duration risk.

⁷The bond pays a coupon: c at $t+1$, ρc in $t+2$, $\rho^2 c$ in $t+3$, and so on.

⁸In general, the term premium can be responsive to the cycle in ECB-BASE: $tp_t = c^{tp} + \rho^{tp} tp_{t-1} + \lambda^{tp} \mathbb{E}_t[OG_{t+40}] + \epsilon_t^{tp}$. The expression $\mathbb{E}_t[OG_{t+40}]$ stands for the expected output gap position of the economy in 10 years and is replaced by a VAR forecast where the future cyclical position is projected based on current economic conditions. In order to avoid a convoluted expectation formation inside the determination of long-term interest rates, the (RE)BASE setting fixes the term premium to its baseline value, $tp = c^{tp}/(1 - \rho^{tp})$. To make comparison viable, we also fix the term premium in the backward-expectation cases for the remainder of the paper.

in Equation (4), except that, in addition to the term premium, a corporate bond spread above risk-free rates is then added to the discounting.

2.1.2 Equity Prices

The description of equity prices follows the Dividend Discount Model. The equity price, P_t^{EQ} , is then equal to the present value of expected future dividend cash flows that are discounted with an appropriate risk factor:⁹

$$P_t^{EQ} = \mathbb{E}_t \sum_{k=0}^{\infty} \frac{c_S D_{t+k}}{(1 + i_t^{COE})^k}, \quad (7)$$

where D_t are dividend payments and c_S is a normalizing constant. Cost of equity, i_t^{COE} , is constructed as $i_t^{COE} = i_t + s_t^{COE}$, with the equity premium, s_t^{COE} , determined as in the original ECB-BASE.

Writing the equity price equation forward, it is cast into a form similar to the bond price equations shown above:

$$P_t^{EQ} = \frac{c_S D_t}{(1 + i_t^{COE})} + \mathbb{E}_t \left[\frac{P_{t+1}^{EQ}}{(1 + i_t^{COE})} \right]. \quad (8)$$

Dividend payouts to shareholders, D_t , can be approximated by dividend income received by the household sector, which is endogenously determined in the model's property income block. The MCE setting therefore features a forward-looking formation of equity prices that encompasses rational expectations on future dividend growth. c_S is calibrated to target a baseline dividend yield of 2.85 percent, consistent with the historical average of the MSCI Dividend Yield indicator.

⁹In the original ECB-BASE, this equation is cast into a backward-looking version of the Dividend Discount Model. As in Fuller and Hsia (1984), it assumes two stages of dividend growth. There is an extraordinary growth phase (characterized by g^{ST}) that lasts for $2H$ years and a stable growth rate phase (g^{LT}) that lasts forever. Future growth of dividends is either projected based on exogenous information or can follow rules that rely on information from the past. Keeping the notation of the main text in all other cases, the backward-looking equity price can be written as $P_t^{EQ} = \frac{c_S D_t}{i_t^{COE} - g_t^{LT}} (1 + g_t^{LT} + H(g_t^{ST} - g_t^{LT}))$.

Finally, net financial assets held by households, FW_t , are affected by changes in stock prices and gain from returns earned on long-term and corporate bonds:

$$\begin{aligned} \frac{FW_t}{FW_{t-1}} = & \omega_0 + \omega_{10Y}(1 + r_t^{10Y}) + \omega_{CB}(1 + r_t^{CB}) \\ & + \omega_{EQ}(P_t^{EQ}/P_{t-1}^{EQ}), \end{aligned} \tag{9}$$

where ω_0 , ω_{10Y} , ω_{CB} , and ω_{EQ} are calibrated shares of assets not subject to revaluation, government debt securities, corporate debt securities, and equities in total financial assets.

2.1.3 Exchange Rate

The nominal exchange rate is determined via an uncovered interest rate parity (UIP) condition that clears the interest rate differential with trading partners. The no-arbitrage condition is that the rate of return on domestic bonds in the domestic currency must equal the rate of return on foreign bonds whose payoff is converted into the domestic currency via the future spot exchange rate:

$$i_t^{10Y} = i_t^{10Y,F} + \mathbb{E}_t(s_{t+1} - s_t), \tag{10}$$

where s_t is the log of the nominal exchange rate of the domestic currency over a basket of trade-weighted currencies, i.e., the nominal effective exchange rate. i_t^{10Y} and $i_t^{10Y,F}$ are the long-term nominal interest rates in the domestic and foreign markets (the U.S.), respectively.

In the original model, the nominal exchange rate is governed by an empirical relationship that tracks the price and real interest rate differentials between the euro area and the foreign market. In order to depart from purchasing power parity, and therefore from a flexible price case, an explicit functional form for the change of the real exchange rate is assumed¹⁰ and constitutes the basis for the estimation. As a result, an (estimated) scaling of interest

¹⁰In BASE, the real exchange rate, defined as $q_t = s_t + p_t^F - p_t$, is assumed to follow $\mathbb{E}_t(q_{t+1} - \bar{q}_{t+1}) = \theta(q_t - \bar{q}_t)$.

rate changes affects the determination of the exchange rate level in ECB-BASE.¹¹

To mimic this setup for the MCE version, we modify the UIP condition such that the expected change of the nominal exchange rate is a scaling of the interest rate differential:

$$\mathbb{E}_t \Delta s_{t+1} = \theta_{uip} \left(i_t^{10Y} - i_t^{10Y,F} \right), \quad (11)$$

where the scaling parameter, θ_{uip} , has been calibrated to match the medium-term dynamics of the exchange rate response to a monetary policy shock of the economy described by the original estimated model version.

In all applications that follow, foreign markets are assumed to be exogenous and therefore do not transmit feedback to the euro-area economy. Any exchange rate response is triggered by domestic interest rate changes. This makes it possible to condition the euro-area economy on an unchanged global environment, thus avoiding modeling complexities which would be beyond the scope of this paper.

2.2 Modified Income

The long-term consumption growth rate, consistent with economic theory, is derived from utility maximization where households consider their lifetime income stream and the value of financial assets (housing wealth and financial wealth). The reader is referred to Angelini et al. (2019) for the full derivation.

The permanent income construct is used to pin down target consumption, $\log C_t^*$. The empirical specification of the target consumption equation, which has been estimated in the original model, is

$$\begin{aligned} \log C_t^* &= \eta_0 + \eta_T \log EY_t^T + \eta_P \log EY_t^P + \eta_D \log W_t^D \\ &\quad + (1 - \eta_T - \eta_P - \eta_D) \log EY_t^L, \end{aligned} \quad (12)$$

¹¹In BASE, the empirical relationship pins down the current period level of the nominal exchange rate directly (not in expectations). The parameter restrictions employed in the estimation imply $s_t = p_t - p_t^F + 1/(1 - \theta) \left(i_t^{10Y} - \bar{\pi}_t^e - (i_t^{10Y,F} - \bar{\pi}_t^{e,F}) \right) + \bar{q}_t$.

where EY_t^i denotes the permanent income type $i \in \{T = transfer, P = property, L = labor\}$, W_t^D denotes observed financial and housing wealth, and η_i stands for propensities to consume out of income type i .¹²

Permanent income is formally derived from the expected present discounted value of future incomes, denoted as J_t :

$$J_t^i = \mathbb{E}_t \sum_{k=0}^{\infty} (1 + r + \phi_0)^{-k} Y_{t+k}^i. \tag{13}$$

Expected permanent income, EY_t , is defined as

$$EY_t^i = (1 - \check{\beta})J_t^i, \tag{14}$$

with $\check{\beta}$ representing the discount factor and equal to $(1 + g)/(1 + \bar{r} - \bar{\pi} + \phi_0)$. $\bar{r} - \bar{\pi}$ is the baseline real interest rate, ϕ_0 is a risk-adjustment factor, and g is the growth rate of potential output.

The expected permanent income of type i can be rewritten in recursive form:¹³

$$EY_t^i = (1 - \check{\beta})Y_t^i + \check{\beta}\mathbb{E}_t(EY_{t+1}^i), \tag{15}$$

where the one-period-ahead expectation of permanent income, $\mathbb{E}_t(EY_{t+1}^i)$, is evaluated as a rational expectation forward term.

2.3 Reformulating the Adjustment Costs

An essential feature of the ECB-BASE model is the significant role played by expectations in the adjustment process and therefore the

¹²We assume that $\eta_T + \eta_P + \eta_D + \eta_L = 1$. Furthermore, although aggregate property wealth can be directly obtained from the data, there are numerous issues related to its measurement. Because of that, we decided to follow FRB/US and postulate that the true property wealth is a weighted average of observed financial and housing wealth and the present value of property income. This results in financial and housing wealth entering the equation as observed and not as the present discounted value of its future streams.

¹³This is different from the original ECB-BASE model where, in order to apply VAR expectations, permanent income is deconstructed into trend and gap terms. The long-run trends are set according to the baseline targets they are supposed to converge to. The gaps to these trends depend on VAR forecasts about the cyclical position of the economy.

stickiness of main macroeconomic variables. The sluggish adjustment of variables to a theory-implied target are modeled through the polynomial adjustment costs (PAC) approach.¹⁴ The benefit of this is that the long-term anchors of the model are consistent with economic theory, but the transition to the target is consistent with empirical relationships. The variables shift toward the long-run targets in an error-correction-type setup. It further captures inertia via a lag structure as well as providing a role for expectations through a forward term.

Because in ECB-BASE, main macroeconomic interactions feature the PAC property, expectations appear in many parts of the model (recall Figure 1 for an overview). For instance, theory suggests that consumption depends (log-)linearly on permanent incomes (log C_t^* in the previous section is the theory-implied consumption target). However, when the model is applied to the data, there is a discrepancy between the observed and the theoretical level and growth rate of consumption. The intertemporal (forward-looking) optimization that determines the transition is modeled via the PAC approach.¹⁵

Now, let y_t be a generic variable of interest and y_t^* be its theoretical counterpart, typically derived from an optimization problem. Assume that the optimizing agents also seek to minimize a cost function of transition, denoted by Θ_t , with respect to the variable of interest, y_t , while taking the target, y_t^* , as given.¹⁶ This cost function is designed to penalize not only deviations between the actual value of the variable, y_t , and the target value, y_t^* , but also changes in the variable itself:

$$\Theta_t = \mathbb{E}_{t-1} \sum_{i=0}^{\infty} \beta^i \left[(y_{t+i} - y_{t+i}^*)^2 + \sum_{k=1}^m b_k ((1-L)^k y_{t+i})^2 \right]. \quad (16)$$

¹⁴For a reference and a more detailed presentation, please see Tinsley (2002).

¹⁵In the case of consumption, the PAC lag structure is akin to internal habit formation in other models. The advantage of the PAC approach, however, is that it is generic and therefore offers flexibility in modeling any macroeconomic variable via this approach.

¹⁶The optimization problem resembles that of an HP filter, yet the instrument employed is distinct.

This cost encompasses the anticipated present value of squared deviations of the decision variable y_t from its target path y_t^* , in addition to adjustment costs associated with the growth rate and higher-order derivatives of the decision variable. Expectations are formed based on the information available up until the end of period $t - 1$.¹⁷ The decision rule derived from the minimization of Θ_t is as follows:

$$\Delta y_t = a_0(y_{t-1}^* - y_{t-1}) + \sum_{k=1}^{m-1} a_k \Delta y_{t-k} + \mathbb{E}_{t-1} \sum_{j=0}^{\infty} d_j \Delta y_{t+j}^*, \quad (17)$$

where the coefficients entering the above decision rule are nonlinear functions of the adjustment cost parameters, b_k , and the discount factor, β . The change in the decision variable depends on deviations from the nonstationary optimal target value, its own lagged values, and the expected infinite discounted sum of future target values. The PAC approach not only introduces persistence to the variable of interest but also creates a dependency on the future values of the target. In contrast to the quadratic cost adjustment optimization problem obtained by setting the number of lags to $m = 1$, as used in Kennan (1979) and Rotemberg (1982), the PAC generalization of the adjustment cost function permits the inclusion of higher-order (frictions) derivatives by setting $m \geq 2$. This will result in a decision rule incorporating an arbitrary number of lags on the endogenous variable and lead changes in the target variable. Note that it is the lagged changes that potentially improve the short-term dynamic properties of the model and its empirical fit.

For ease of exposition, now define auxiliary variable Z_t , which captures the present value forward term in the PAC approach:

$$Z_t = \mathbb{E}_{t-1} \sum_{j=0}^{\infty} d_j \Delta y_{t+j}^*. \quad (18)$$

¹⁷Notice that when considering the PAC framework, the expectations are formed by using the information available at time $t - 1$ and not at t as is was the case in the previous sections. Although it might seem that the issue of including lagged versus current information when forming expectations in the PAC-based model blocks might be somewhat inconsistent, there is no clear-cut argument promoting either. By the fact that all right-hand side variables of Equation (16), and subsequently of Equation (17), are formed based on the same filtration, inconsistency is not present (Laubach and Reifschneider 2003).

Under MCE, the expected path of the target variable can be directly inferred from the system-wide solution of the model.¹⁸ This is contrary to the VAR case. While the VAR, in principle, may be considered a sufficient statistic representation of the full model, expectations are still formed with a restricted information set. In certain applications, such as the analysis of permanent policy change effects, this may result in persistent expectation errors under the VAR setup (see, for example, Brayton, Davis, and Tulip 2000).

Now, the solution under MCE requires rewriting the infinite sum in Equation (18) into a recursive finite-lead representation in terms of Z_t and Δy_t^* , a discount factor, β , and the polynomial coefficients, α_i , which are recoverable from the estimated coefficients, a_i .¹⁹ After some algebraic manipulations, one can derive the MCE representation of the PAC expectation term:

$$Z_t = - \sum_{i=1}^m \alpha_i \beta^i Z_{t+i} + a_0 \left(\Delta y_t^* - \sum_{k=1}^{m-1} \sum_{j=k}^{m-1} \alpha_{j+1} \beta^{j+1} \Delta y_{t+k}^* \right). \quad (19)$$

Note that this setup allows us to retain the estimated parameters, α , which have been estimated under the VAR expectation assumption. It also allows for convenient switching between VAR expectations and MCE depending on how Z_t is determined. Since the adjustment costs feature in many model blocks, the researcher may conveniently decide between hybrid expectation setups, where some sector adjustments can follow VAR-based expectations and others MCE.

2.4 *Dynare's Computational Environment for ECB-(RE)BASE*

All the model's simulations and estimates are carried out using Dynare's new features. Dynare provides commands to define the

¹⁸Under VAR-based expectations in the original model, Z_t is expressed as $Z_t = \sum_{j=0}^{\infty} d_j \iota H^{j+1} x_{t-1} = h' x_{t-1}$, where coefficient vector h' is a nonlinear function of VAR matrix H , discount factor β and a_k coefficients. x are the VAR variables which are a subset of the full model variable set. ι is a selection vector that picks variables in x that are considered to be relevant for the imperfect information set under which the VAR forecast is formed.

¹⁹With $\alpha_{m-1} = a_{m-1}$, $\alpha_i = -\Delta a_i$ for $i = 2, \dots, m-2$ and $\alpha_1 = a_0 - a_1 - 1$.

backward auxiliary model used to form expectations, `var_model` and `trend_component_model`,²⁰ by targeting a set of equations in the `model` block. Any expected variable in the model can be computed using the auxiliary model, provided that the variable to be expected is part of the auxiliary model and that the expectation operator is applied directly to the variable.²¹ This is done by defining an expectation model with the `var_expectation_model` command. It refers to a previously defined auxiliary model and allows setting of the anticipated variable and the expectation horizon. Then, in the `model` block we can use the `var_expectation` keyword, which must be linked to an expectation model, as a replacement for a forward variable. We use this approach for all expectations, including those not related to PAC equations.

Dynare also provides commands related to the target and the anticipations in the PAC equation. In the simplest situation, where the target of the PAC equation is a single nonstationary variable, y^* in Equation (17), Dynare automatically identifies the target by parsing the error-correction term. This is the case for all the ECB-(RE)BASE PAC equations, but composite targets are also allowed:²² the target is then a linear combination of nonstationary and stationary variables. The command `pac_model` defines the model for a PAC equation. For VAR-based expectations, the PAC model must be linked to a previously defined auxiliary model. If this reference is not given, Dynare understands that the expectations must be model-consistent (MCE), as in Equation (19). This command is also used to set the discount factor and, if necessary, defines the correction for growth neutrality. Once the PAC model has been defined, we can substitute the operator `pac_expectation` for Z_t in the notation above. Dynare replaces it with an auxiliary variable, whose definition can be either Equation (19) or the VAR forecast, depending on whether there is a link to an auxiliary model in the PAC model, and computes the necessary parameters (parameters α_k or vector h).

²⁰See the reference manual, Adjemian, Juillard et al. (2024), for a full description of these commands and those introduced below. Note that, except for the PAC equation, all the estimation routines are undocumented.

²¹The latter condition differs from what is assumed by default in Dynare, where the conditional expectation applies to the entire (nonlinear) equation.

²²See the commands `pac_target_info` and `pac_target_nonstationary` in Adjemian, Juillard et al. (2024).

The PAC equations are estimated under VAR-based forecasts. The estimation of the autoregressive parameters, a_k ($k = 0, \dots, m$), is not obvious since the vector h in the VAR depends nonlinearly on the autoregressive parameters. In the literature, this equation is typically estimated through iterated ordinary least squares (as in FRB/US). However, Dynare also permits estimation of this equation using nonlinear least squares (NLS), which is the favored approach here.

The simulation of the model uses simulation routines specialized either for backward models if all the expectations are VAR-based, or relying on the perfect foresight solvers or extended path simulation routine if some expectations are model-consistent. Note that Dynare can reduce computation time by exploiting the property that each equation of the semi-structural model is written as an endogenous variable equal to an expression.

3. Monetary Policy Transmission across Alternative Expectations

This section explores how transmission dynamics change by incorporating model-consistent expectations in the model economy. We focus on investigating the propagation of monetary policy shocks, as they are a major driver of economic variation and are of key interest for the central bank's model-based analysis.

All simulations are conducted relative to a baseline. Initially, we establish the model's balanced-growth path (BGP) by executing a sufficiently long simulation of the respective version of the model, ensuring that the system settles at constant growth rates. Consequently, the growth rates of stationary variables become zero, while those of nonstationary variables align with their theoretical long-run growth rates. The initial simulation periods, where the system has not yet stabilized, are discarded. Thereafter, any further simulations can be performed around this stable BGP.

3.1 Alternative Expectations Settings

In the following analysis, we implement a shock of identical magnitude in the model economy while sequentially switching the expectation settings of various sectors to the MCE version. Table 1 provides

Table 1. The Role of Expectations in the Model's Blocks

	VAR Expectations	Financial MCE	Financial + PCs MCE	Financial + PCs + Income MCE	Full MCE
Exchange Rate/UIP		✓	✓	✓	✓
Equity Prices		✓	✓	✓	✓
Bond Prices/Interest Rates		✓	✓	✓	✓
Wage Setting			✓	✓	✓
Price Setting*			✓	✓	✓
Permanent Income				✓	✓
Household Consumption PAC					✓
Business Investment PAC					✓
Residential Investment PAC					✓
Employment PAC					✓
House Prices PAC					✓
Property Income PAC					✓

* The medium-term inflation target of the central bank remains imperfectly anchored throughout this paper. It is calibrated to be of low sensitivity with respect to current price inflation and therefore not influence simulations noticeably.

Note: The first column refers to model blocks where expectations play a role in the BASE model. The columns describe the expectation modalities considered when switching expectations. A check mark means that the VAR expectations/the backward rule of the equation in the BASE model is switched to its MCE version of (RE)/BASE. PCs = Phillips curves.

an overview of the modalities we consider. Between the pure VAR expectations setting and the “Full MCE” setting, there exists a spectrum of hybrid expectation settings. In the context of examining the effects of monetary policy, these hybrids may prove useful, as they reflect the degree to which certain sectors understand policy signals. For instance, Bernanke, Kiley, and Roberts (2019) use expectation settings where the private sector is backward-looking, serving as a stand-in for imperfect credibility of policy shifts. Conversely, simulations under the Full MCE setting, i.e., rational expectations, are related to full credibility.

First, we consider the scenario where only the financial sector is forward-looking. As argued by Bernanke, Kiley, and Roberts (2019), financial markets may be perceived as having greater incentives, expertise, and resources to comprehend and evaluate the implications of monetary policy. This setting is particularly relevant when policy shifts are considered temporary, requiring significant effort to discern their consequences, while households and nonfinancial firms may choose to fully invest resources only in the case of more permanent changes.

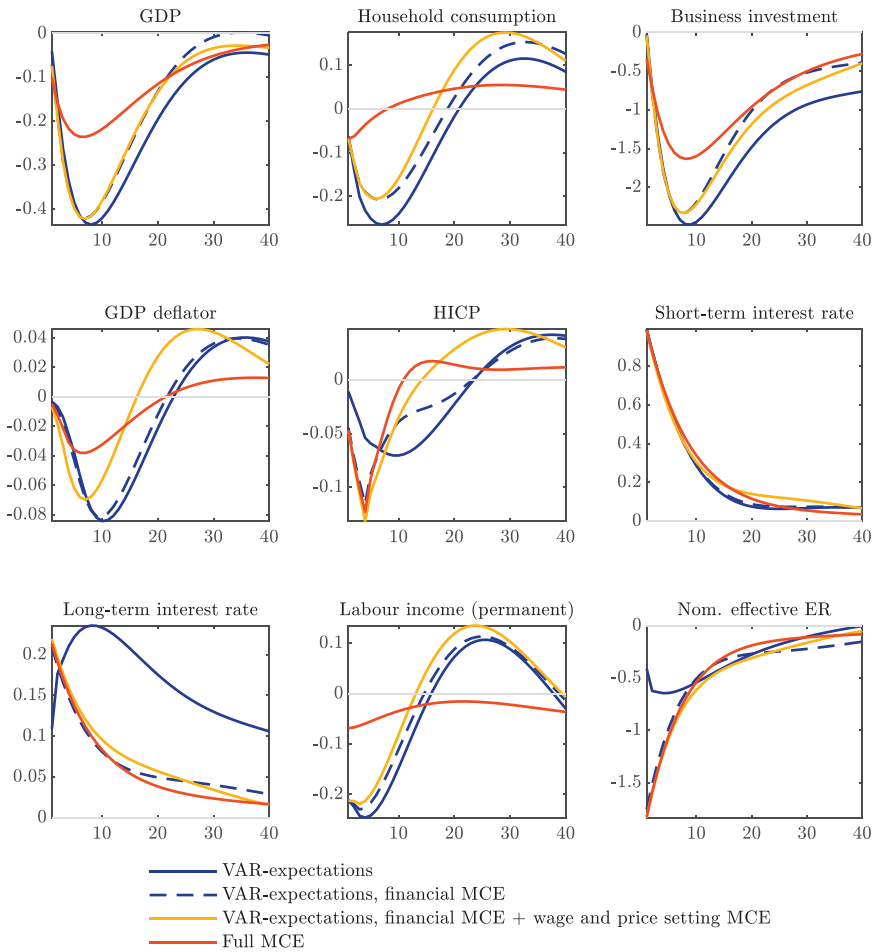
Then, we consider the supply side, and therefore the nominal sphere, as forward-looking. This can be motivated by goods-producing firms possessing comparatively more means and incentives to understand policy changes than households. In this sequence, households and therefore consumers are least likely to be forward-looking. Switching income evaluations to MCE can be thought of as extending rational expectations to this sector as well. Finally, all transmission dynamics (PAC) are adjusted to be forward-looking in the Full MCE setting.

Sectors adhering to VAR expectations form future expectations based on past regularities. This may serve as a reasonable approximation of the true model during normal times, but the Full MCE is likely to yield significantly different insights when analyzing shifts during specific crises or nonlinearities, such as the Lower Bound period on interest rates (Reifschneider and Roberts 2006).

3.2 Simulations of Policy Shocks

Figure 2 presents impulse response functions for a standard monetary policy shock, sequentially swapping expectations from the

Figure 2. Monetary Policy Shock (100 Basis Points)—Switching Expectations in Sectors



Note: Responses to an unanticipated increase in the short-term nominal interest rate by 100 basis points. All simulations are conditioned on the term premium and financing spreads being fixed to their baseline values. The horizontal axis shows time in quarters. The vertical axis depicts short- and long-term interest rates as annualized percentage point deviation from the baseline value. GDP deflator inflation and HICP inflation are shown in percentage point deviation from the baseline year-on-year growth rate. All other variables are in percent deviation from their baseline levels.

VAR-based version to the Full MCE version (see the channels that are switched on in the overview in Table 1). The economy responds to an unanticipated shock to the monetary policy rule of the model, such that the short-term interest rate increases by 100 basis points upon impact.

The ECB-BASE VAR expectation response is depicted by the solid blue line. It is important to note that the increase in the monetary policy rate has expectational features in this backward-looking setting also. The short-term interest rate is incorporated into the core VAR of ECB-BASE, thereby influencing all VAR-based expectations. Higher short-term interest rates permeate through the yield curve since the risk-free long-term rate depends on expected future short-term rates. Alongside an endogenous rise in spreads, this shift in interest rates elevates borrowing costs for households and firms, consequently dampening consumption expenditures as well as business and residential investment. The reduction in aggregate demand exerts downward pressure on nominal wages and prices via the Phillips curve mechanism. Additionally, monetary policy tightening leads to an appreciation of the exchange rate, and the resultant loss in export market shares further exacerbates the contractionary effect on economic activity.

Next, we switch the financial block of the model to MCE (represented by the dashed blue line). This results in a more abrupt repricing in the yield curve, which gradually unwinds, contrary to the hump-shaped response suggested by VAR-based expectations. This directly translates into a more moderate reaction in financing costs, which is a central propagation channel and results in a lesser decline in investment and consumption expenditures. Moreover, the revaluation of nominal assets is less negative, thereby exerting less downward pressure on consumption decisions compared to the VAR-based expectations case. In line with the dampened demand cycle, the GDP deflator also exhibits a more subdued response. The determination of the exchange rate changes with financial MCE causes the euro to appreciate in a more front-loaded fashion, which directly impacts HICP inflation. This is due to an exchange rate repricing channel affecting imported energy, which will be discussed in more detail later. But overall, switching the financial sector to a forward-looking setting has limited implications on the real side. This is because a milder pass-through to financing rates under the

financial MCE specification is offset by a stronger euro appreciation, triggering a worsening of the trade position.

We proceed by switching the price and wage Phillips curves to the MCE mode (illustrated by the yellow line). This model version is to be understood as an economy where firms are strongly forward-looking, in addition to the financial sector, but households are not. The domestic price response is somewhat front-loaded compared to the VAR expectation case, albeit not amplified. Since the wage setting also becomes more front-loaded, income and, consequently, consumption dynamics experience some changes. Other real expenditures, however, do not exhibit substantial changes.

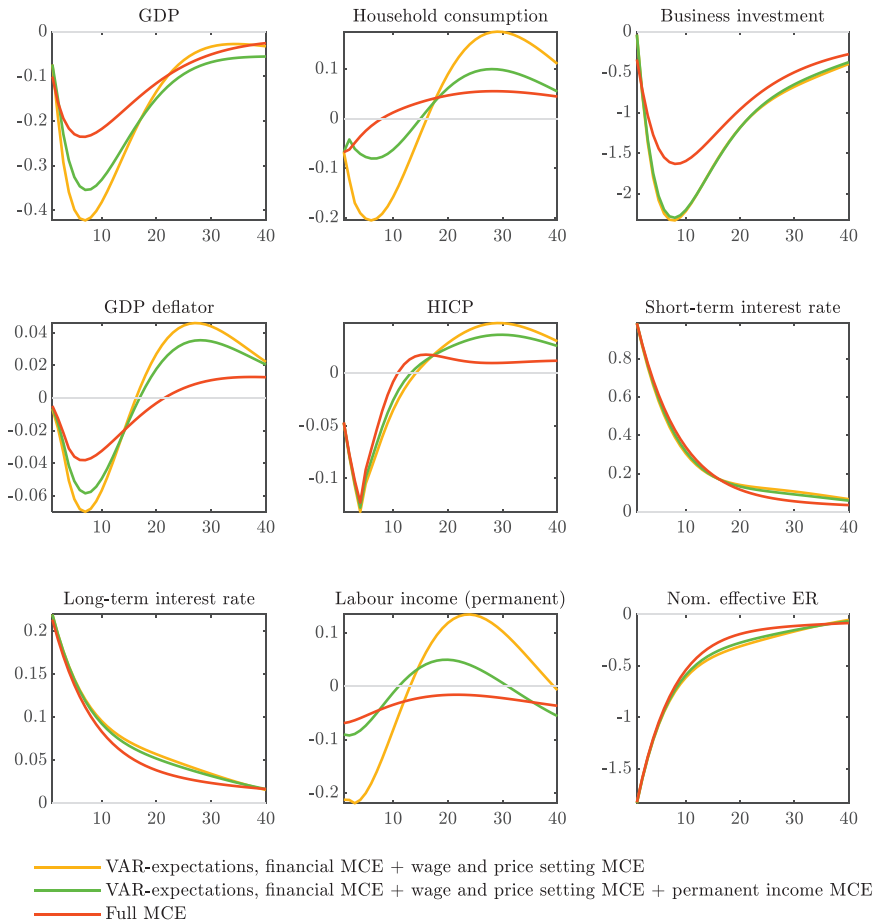
Finally, we switch the entire model into Full MCE mode (depicted by the red line). The transmission of the monetary policy shock becomes smoother than in the previous settings. A significant factor is the permanent income perception, which strongly influences consumption decisions. The household sector now anticipates the full future repercussions of the cycle. As a result, consumers tend to “look through” the demand deterioration and its implications for their lifetime income. Consumption, and thus economic activity and other expenditures, are smoothed out compared to the more backward-looking expectation settings. Moreover, with all adjustment dynamics now operating in forward mode, transitions become more gradual, leading to considerable smoothing of both expenditure and output.

3.2.1 Permanent Income Smoothing

In the (RE)BASE model, model-consistent expectations do not amplify the effect of unexpected standard monetary policy shocks on output. Our next step is to disentangle the influence of income smoothing from the effect of smoother transition dynamics that are due to switching PAC expectations. This analysis is particularly useful as it highlights a unique feature of our model: unlike typical DSGE model specifications, consumption behavior is not driven by an Euler equation where the real interest rate directly steers consumption-saving decisions. Instead, our model focuses on empirical relationships concerning the propensity to consume from various sources of income and wealth.

To explore the significance of the permanent income channel on its own, Figure 3 illustrates another model version adding to the

**Figure 3. Monetary Policy Shock (100 Basis Points)—
The Role of Income Smoothing**



Note: Responses to an unanticipated increase in the short-term nominal interest rate by 100 basis points. All simulations are conditioned on the term premium and financing spreads being fixed to their baseline values. The horizontal axis shows time in quarters. The vertical axis depicts short- and long-term interest rates as annualized percentage point deviation from the baseline value. GDP deflator inflation and HICP inflation are shown in percentage point deviation from the baseline year-on-year growth rate. All other variables are in percent deviation from their baseline levels.

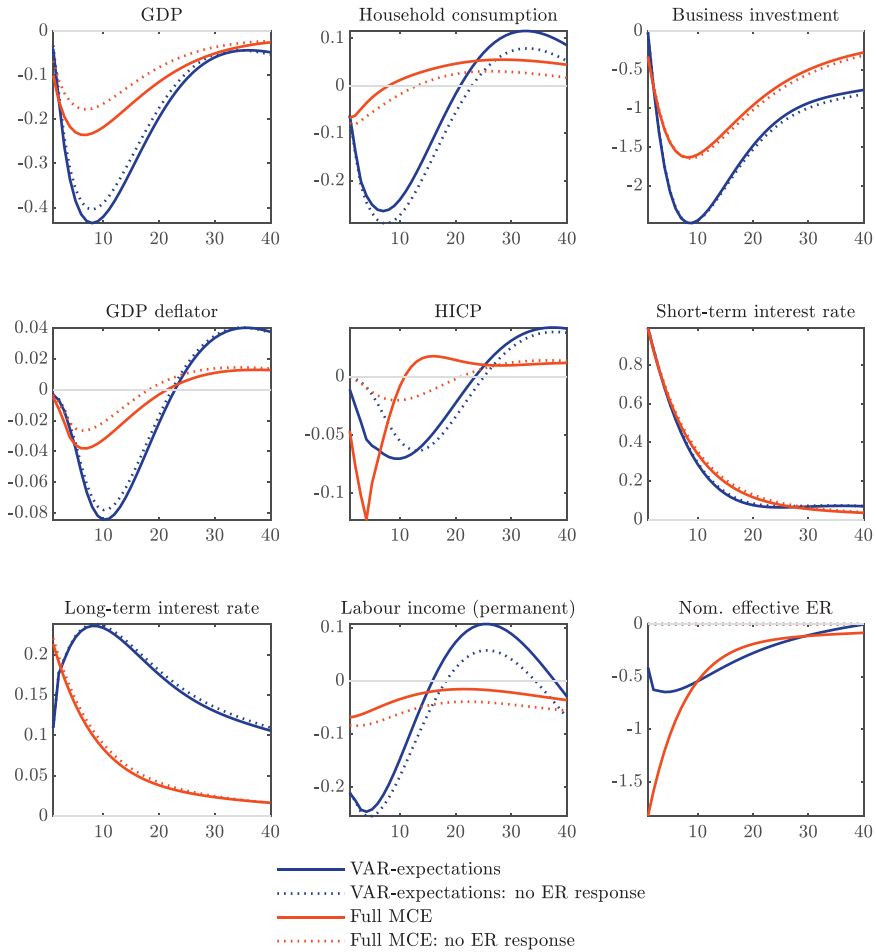
previous exercise. We begin with the hybrid setting where the supply side has been switched to MCE. Then, represented by the green line, the permanent income perception only is changed to MCE. This is done without yet considering MCE in the PAC setting. The permanent income MCE accounts for more than half of the cyclical change in consumption when comparing the hybrid modality to the Full MCE modality. However, investment is not affected to the same extent. Since consumption is smoothed, the decreases in output and prices are lower—proportional to the share of consumption. The remaining difference compared to the Full MCE is attributed to transitioning the PAC specifications to a model-consistent expectation mode.

3.2.2 The Role of the Exchange Rate

The analysis thus far has been conducted by considering a response of the exchange rate to the interest rate differential between domestic and foreign (U.S.) interest rates. The exchange rate channel is a crucial aspect of monetary policy transmission and merits separate investigation. The deflators for import and export prices are modeled based on domestic and external cost pressures, along with considering the role of price competition in international trade. Consequently, the exchange rate directly influences the prices at which euro-area exports compete in third markets. Additionally, the model explicitly accounts for the price of oil and gas energy. The euro area is heavily reliant on energy imports, which are predominantly priced in U.S. dollars. Therefore, the EUR/USD exchange rate impacts GDP deflator inflation by affecting the price of the inputs to production. The model also entails a link of HICP energy inflation to these international market prices when converted to domestic prices. Therefore, the model specifically accounts for the prominent role of household expenditures on petrol and heating.

Figure 4 begins with the VAR expectation case, where the financial sector is in MCE mode as a backward-looking benchmark version. It is plotted against the Full MCE case. There is no feedback of the global environment to changes in the euro area. An expansionary monetary policy in the euro area subsequently triggers an appreciation of the euro as domestic interest rates rise and opens the differential with unchanged foreign markets. The exchange rate

**Figure 4. Monetary Policy Shock (100 Basis Points)—
The Role of the Exchange Rate Channel**



Note: Responses to an unanticipated increase in the short-term nominal interest rate by 100 basis points. All simulations are conditioned on the term premium and financing spreads being fixed to their baseline values. The horizontal axis shows time in quarters. The vertical axis depicts short- and long-term interest rates as annualized percentage point deviation from the baseline value. GDP deflator inflation and HICP inflation are shown in percentage point deviation from the baseline year-on-year growth rate. All other variables are in percent deviation from their baseline levels.

response in the backward-looking version is notably more gradual, while the Full MCE case displays a more immediate loading with a stronger mean reversion. This is a known feature in the literature, where the DSGE approach to specifying a UIP condition on expected exchange rate changes typically does not lead to enough persistence to generate a hump-shaped response (Adolfson et al. 2008).

Illustrated by lightly dotted lines, the figure demonstrates alternative economic responses under an unchanged exchange rate. This switches off the loss of export market shares and the deterioration in net trade. Output therefore declines by less. It also switches off the endogenous repricing of energy imports quoted in domestic currency. In the endogenous model, an appreciation results in a real income gain because energy prices become cheaper. This channel mitigates the negative impact of production loss due to lower exports but does not fully compensate for it.

The significance of the exchange rate channel for output dynamics is slightly stronger under Full MCE. Given that HICP energy inflation is a significant contributor to total HICP inflation, the exchange rate influences consumer prices more directly, as opposed to the more gradual pass-through to producers' factory gate prices. Without the exchange rate channel, the differences in HICP dynamics across expectations are more akin to the dynamics of GDP deflator inflation. Overall, this channel somewhat obscures the heterogeneity of the importance of expectation settings for price developments.

3.2.3 Anticipation of News Shocks

A main capability of our forward-looking expectation model is its ability to analyze anticipated news shocks. It has been shown that these news shocks, in addition to surprise shocks, play a quantitatively significant role in business cycle fluctuations (Schmitt-Grohé and Uribe 2012). It has been demonstrated that theoretical models can be effectively augmented with news shocks to generate impulse responses that exhibit comovements consistent with empirical data (Barsky and Sims 2011).

In (RE)BASE, a news shock on the monetary policy rate is implemented as follows. The nominal short-term interest rate reacts as

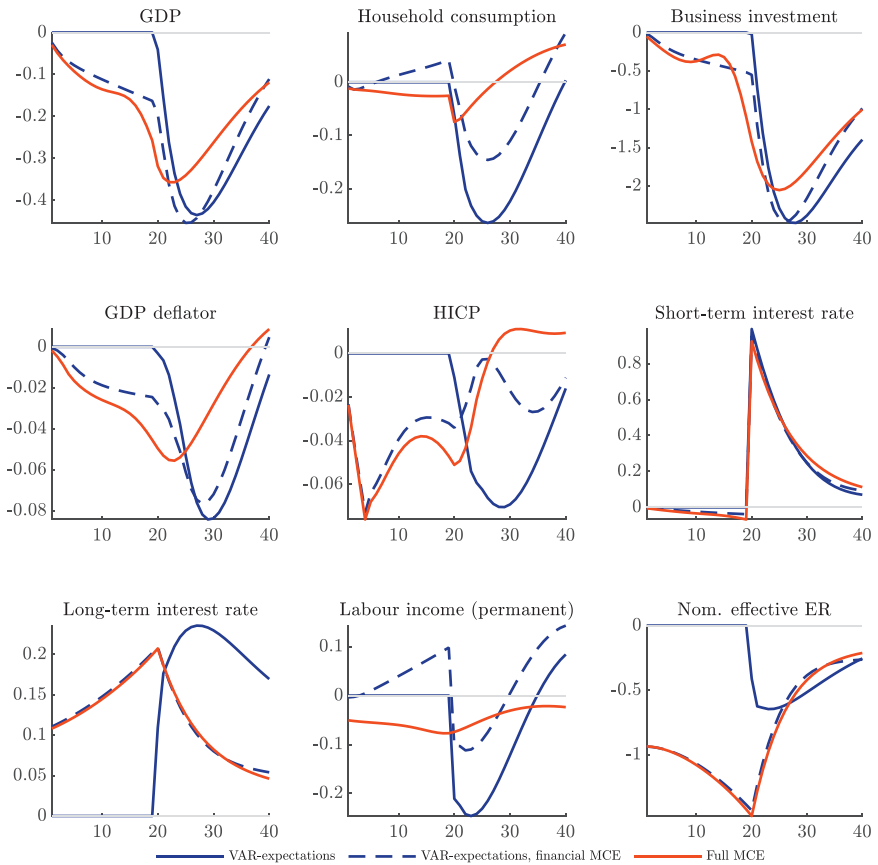
$$i_t = \beta_1 i_{t-1} + (1 - \beta_1) [\bar{i} + \beta_2 \hat{\pi}_t] + \beta_3 \Delta \pi_t + \beta_4 \Delta y_t + \epsilon_{0,t}^i + \epsilon_{1,t-1}^i + \dots \quad (20)$$

The first part of the equation is the model's Taylor rule, reflecting the interest rate's response to the inflation gap, $\hat{\pi}_t$, the change in the inflation rate, $\Delta\pi_t$, and output growth, Δy_t . As the simulations are conducted under perfect foresight, the model economy resolves under a known path of shocks, $\epsilon_{k,t-k}^i$. It is important to note that this process is independent of how the private sector forms expectations. Technically, anticipated news shocks can be simulated across all different expectation settings. An economy relying on backward-looking information would not yet respond to a news shock related to the future. Under rational expectations in (RE)BASE, however, this information influences the mechanics as soon as the shock path becomes known, referred to as the date of announcement.

In the spectrum of credibility discussed above, this can be seen as two polar cases. If the model's expectations are fully VAR-based, then a future policy change cannot be credibly communicated, surprising the economy if it materializes. Conversely, Full MCE implies full anticipation under perfect foresight, connoting full credibility of central bank communication.

Figure 5 illustrates the model economy's responses to an announced increase in the short-term interest rate, 20 periods ahead. Under VAR expectations, the contractionary shock is a complete surprise in period 20. The model's response function is equivalent to shifting an unanticipated shock to period 20. Rational expectations active in the financial block then serve as another benchmark case. It reflects a scenario where only the financial sector understands the monetary policy shift while other sectors do not, emulating an imperfect credibility setup. The long-term interest rate reprices upon announcement at $t = 1$, leading to an immediate increase in firms' and households' borrowing costs. Investment and consumption, and thus GDP, begin to decrease even before the actual hike in the monetary policy rate. Similarly, the value of nominal financial assets also declines immediately, further weighing on consumption. Importantly, the exchange rate also reprices at the time of announcement, adding another layer of anticipation. The appreciation fuels the decrease in output through declining exports, though it also leads to a rapid decline in HICP energy prices, amplifying the response of total HICP.

Figure 5. Monetary Policy Shock (100 Basis Points)—Anticipated News Shock



Note: Responses to an announced increase in the short-term nominal interest rate by 100 basis points, 20 periods in the future, across expectation settings. All simulations are conditioned on the term premium and financing spreads being fixed to their baseline values. The horizontal axis shows time in quarters. The vertical axis depicts short- and long-term interest rates as annualized percentage point deviation from the baseline value. GDP deflator inflation and HICP inflation are shown in percentage point deviation from the baseline year-on-year growth rate. All other variables are in percent deviation from their baseline levels.

In the Full MCE model, considering a fully credible policy announcement leads to price and wage dynamics that are somewhat front-loaded compared to the benchmark case. Notably, households' evaluation of permanent income becomes much smoother,

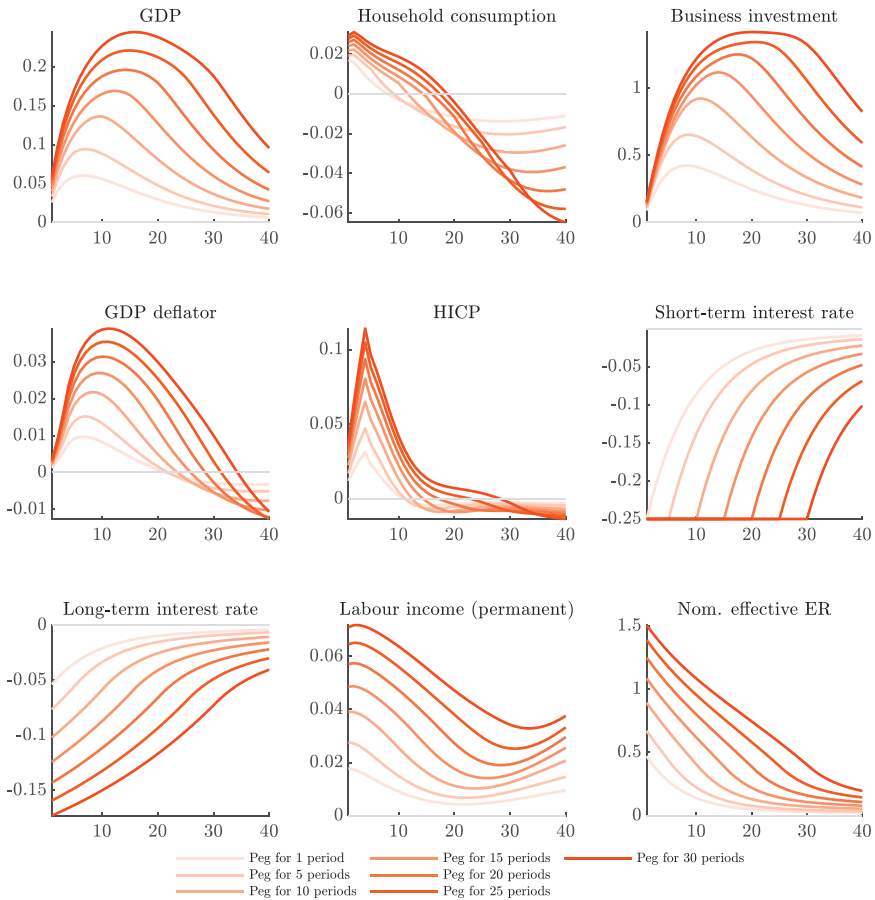
with a significant decrease already during the anticipation phase. This results in consumption smoothing over the full cycle, with a stronger decrease up front but more positive dynamics after the short-term interest rate increase actually becomes apparent at $t = 20$.

3.2.4 Forward Guidance

In the previous exercise, the Taylor rule is active throughout, albeit only moderately leaning against the cyclical downturn during the anticipation phase. The steady interest rate policy, however, does not create a forward guidance puzzle in ECB-(RE)BASE. Many New Keynesian models face the phenomenon of generating explosive economic dynamics at a fixed interest rate for prolonged periods, an issue that has been of focal attention during times when the interest rate was at the Zero Lower Bound (see, for example, Cochrane 2017). Recently, Maliar and Taylor (2024) demonstrated that the puzzle arises at literally fixed rates as well as for policy rate rules that are insufficiently responsive to contain future shocks.

To systematically investigate the role of forward guidance in our model, we simulate exercises similar to those in Christoffel et al. (2022). The authors show that without mitigating the puzzle, a continuous peg of the interest rate for 10 periods creates explosive economic behavior in the New Area-Wide model, a medium-scale DSGE model of the euro area. Figure 6 presents exercises conducted with (RE)BASE, where the central bank pegs the interest rate to a value 25 basis points below the baseline value for different durations. The economic dynamics shift proportionally to the length of the peg and do not become explosive. A major difference from other New Keynesian models is the absence of intertemporal reallocation forces on consumption implied by an Euler equation. Christoffel et al. (2022) tamed the forward guidance puzzle in the DSGE by introducing additional discounting and conditioning expectations on survey data. Both factors limit the importance of future stimulus implied by holding the policy rate fixed for an additional quarter. The parametrization of ECB-(RE)BASE, as a strongly data-conditioned forecast model, already displays relative dynamic stability without further interventions.

Figure 6. Forward Guidance—Peg of the Policy Rate (–25 Basis Points) for Multiple Periods



Note: Responses to an anticipated and credible peg of the short-term nominal interest rate by 25 basis points below the baseline value for multiple periods. Expectations modality: Full MCE. All simulations are conditioned on the term premium and financing spreads being fixed to their baseline values. The horizontal axis shows time in quarters. The vertical axis depicts short- and long-term interest rates as annualized percentage point deviation from the baseline value. GDP deflator inflation and HICP inflation are shown in percentage point deviation from the baseline year-on-year growth rate. All other variables are in percent deviation from their baseline levels.

4. Recent Supply Shocks through the Lens of Alternative Expectations

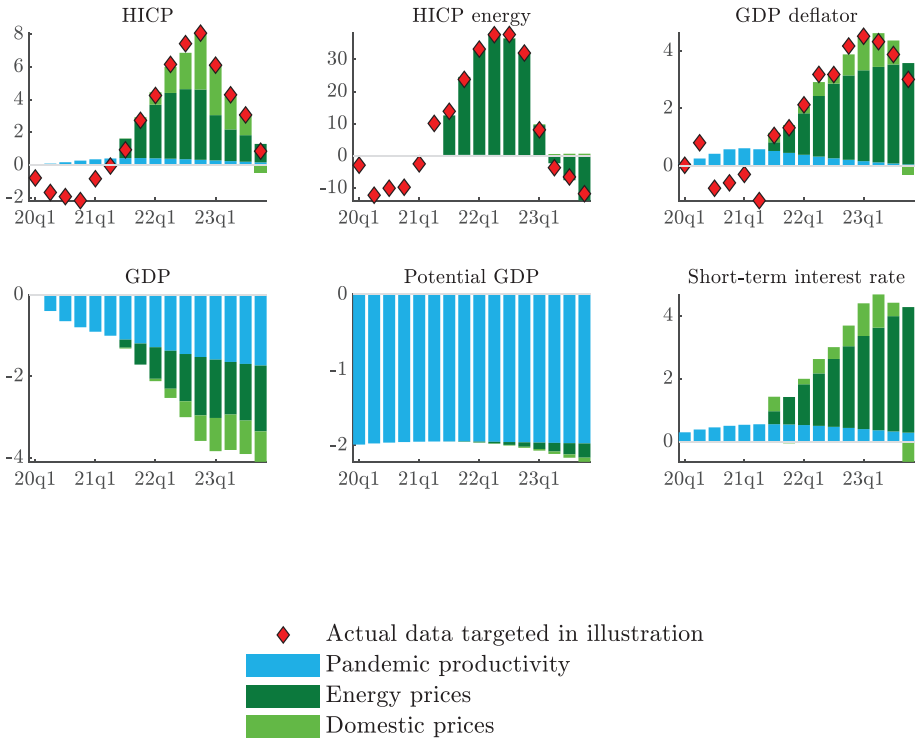
In this section, we apply the model to a real-world analysis. Recent supply shocks have exerted extraordinary inflation pressures on the euro-area economy. The resulting price movements were significant enough to provoke notable changes in the long-term inflation expectations of the private sector. This situation created a trade-off for central banks: the desire to “look through” these supply shocks to avoid exacerbating the cyclical downturn with interest rate hikes, versus the necessity to reanchor inflation expectations by demonstrating a firm commitment to combating above-target price pressures as they emerge. In terms of economic model interpretation, it has been shown that a full forward-looking setting implies an easier reanchoring of inflation expectations (Beaudry, Carter, and Lahiri 2023). Using (RE)BASE, we aim to explore how different expectations frameworks might lead to varying assessments of the trade-off between price stability and output that the central bank faces.

4.1 Scenario Design: Shock Calibrations

We construct a supply-shock scenario that is reminiscent of key characteristics of the recent inflation surge period. The scenario focuses solely on supply shocks. It is designed as follows: we activate three types of residuals to create a supply-shock scenario. Their contributions to the price increases and resulting macroeconomic dynamics through the model (considering VAR expectations) are documented in Figure 7.

First, we employ a stylized productivity shock that reduces potential GDP by approximately 2 percent from the onset of the COVID-19 pandemic. The shock aims to illustrate a contribution of the pandemic-induced supply-side pressures on prices. That may entail production and delivery bottlenecks as well as legacy effects from the lockdowns and containment policies on capacity. In order to calibrate this shock, we refer to the findings of Angelini et al. (2020), who introduced a pandemic module into ECB-BASE in order to study the interaction of the macroeconomy when individuals are exposed to, infected by, and recover from (SIR) the COVID-19 virus. According to this study, around 45 percent of the cumulative output loss during the pandemic is reflected into potential output. In our

Figure 7. Construction of a Supply-Shock Scenario—VAR Expectations



Note: The pandemic productivity shock is simulated as of 2020:Q1, while the other shocks are assumed to arise as of 2021:Q3. HICP, HICP energy, and GDP deflator inflation are shown in percentage point deviation from the year-on-year growth in the baseline. The data are shown as the deviation from a 2 percentage point baseline growth rate. The short-term interest rate is shown as the annualized percentage point deviation from the baseline value. All other variables are in percent deviation from their baseline level values.

exercise, we illustratively assume a 4 percent decrease in the level of output persisting even after the pandemic health crisis subsides.

Second, the scenario incorporates energy price shocks. These shocks are modeled by adjusting the path of imported oil prices to match the observed increase in HICP energy inflation in the euro area starting as of 2021:Q1.²³ This profile is matched until the end

²³See Table 2 for an overview of the data targeted in this scenario.

Table 2. Data Sources in the Construction of the Scenario

Observable	Raw Metric	SDW Codes	Further Treatment
HICP Inflation	Annual rate of change	ICP.M.U2.N.000000.4.ANR	Seasonally adjusted, quarterly average
HICP Energy Inflation	Annual rate of change	ICP.M.U2.N.NRGY00.4.ANR	Seasonally adjusted, quarterly average
GDP Deflator Inflation	Index	MNA.Q.Y.I9.W2.S1.B.B1GQ. Z. Z. Z.IX.D.N	Annual rate of change
Source: SDW (ECB Statistical Data warehouse).			

of our simulation horizon, 2023:Q4. The oil price mechanism in the model also entails a pass-through to the GDP deflator, capturing the effects of imported carbon energy sources used as inputs to production. Since our scenario targets HICP energy prices, we effectively capture a general imported energy price increase, which was largely due to elevated gas prices during this episode. The oil price shock mainly transmits as a loss of real income and, consequently, as a loss in purchasing power, leading to decreased consumption, production, and other expenditures. The central bank adjusts interest rate policies to stabilize the gap between GDP deflator growth and the target rate.

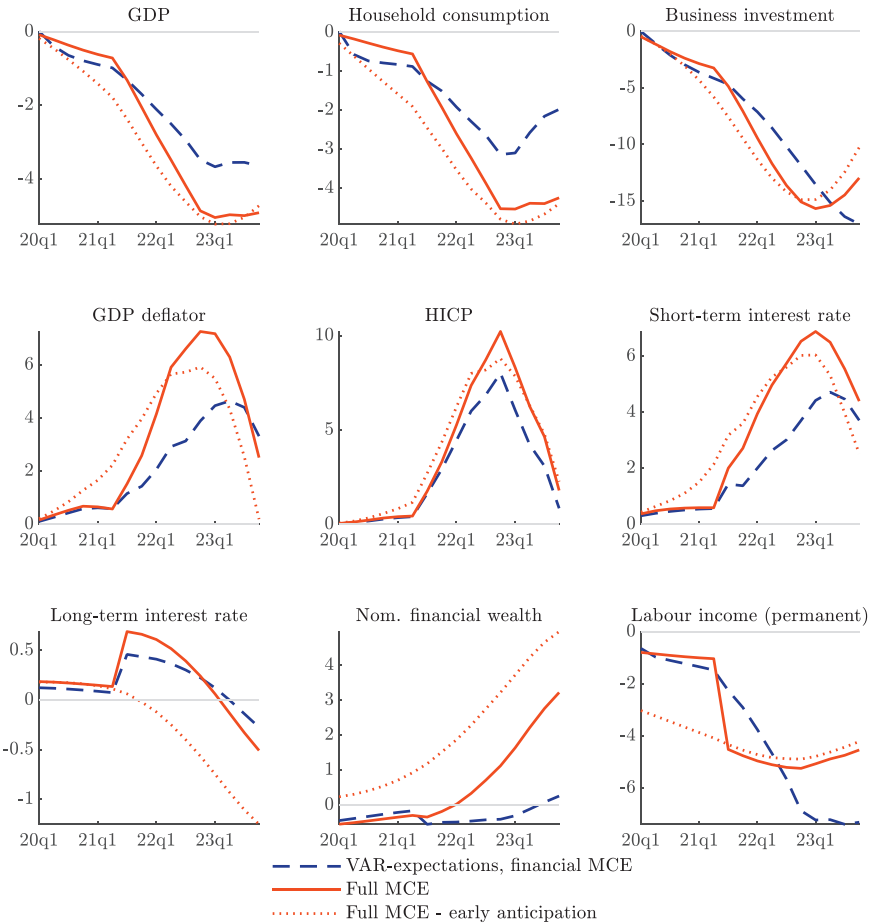
While the productivity shock contributes to the inflation surge scenario to a limited extent, adding the energy price shock would result in an increase in GDP deflator inflation comparable to what was observed in actual data. Overall, this initial combination would account for approximately half of the HICP inflation increase above the central bank target rate noted from 2021:Q3 to 2023:Q4. To complete the scenario, we introduce a shock that matches the remaining GDP deflator and HICP inflation deviations in the data through markup shocks. The markup shocks directly shift the GDP deflator (changing the residual of the main Phillips curve pricing equation) and the HICP excluding energy and food inflation. These shocks further contribute to price increases, triggering additional losses in real disposable income and leading to a further decline in consumption.

4.2 Supply Shocks across Expectation Settings

In this section, we keep the sequence of shocks fixed, but resimulate the model under backward expectation with the financial sector being forward-looking and under Full MCE. Figure 8 displays the main simulations as dashed blue and solid red lines. Each simulation displayed is created by the same combination of shock sequences constructed above. The interested reader is referred to the working paper, Adjemian, Bokan et al. (2024), for a systematic documentation of switching the expectations for each shock separately.

We begin with a backward-expectations setting, where only the financial sector is forward-looking and forms rational expectations. We assume an endogenous monetary policy where the nominal short-term interest rate starts to rise at the first signs of price

Figure 8. Macroeconomic Dynamics in the Supply-Shock Scenario across Expectation Modalities



Note: Responses to the same sequence of shocks, as constructed in Figure 7, across alternative expectation settings. In the first two simulations, the energy and domestic markup shocks are known to the economy as of 2021:Q3. In the simulations labeled “early anticipation,” the energy and domestic markup shocks are anticipated as of 2020:Q1 even if they materialize in 2021:Q3. All simulations are conditioned on the exchange rate, the term premium, and financing spreads being fixed to their baseline values. The graph depicts short- and long-term interest rates as the annualized percentage point deviation from the baseline value. GDP deflator inflation and HICP inflation are shown in percentage point deviation from the baseline year-on-year growth rate. All other variables are in percent deviation from their baseline levels.

increases from the pandemic leg of the scenario. The other shocks regarding energy and domestic price markups unexpectedly impact the economy as of 2021:Q3. The short-term interest rate then adjusts accordingly, and spills into the yield curve. Due to anticipation in the financial sector, the long-term interest rates reflect the expected repricing of the short-term rate path at the beginning of the markup shock sequence. The private sector gradually perceives the real income loss implied by the price increase, causing the consumption path to adjust correspondingly.

When simulating the inflation surge scenario under Full MCE, the markup shocks lead to a significantly larger price increase because the Phillips curve expectations front-load the nominal implications. With endogenous monetary policy, this would necessitate a quicker and more pronounced rise in interest rates. Furthermore, households anticipate and front-load the perceived real income loss. Overall, households sharply reduce their consumption under the Full MCE setting, resulting in a more pronounced decline in GDP compared to the backward expectations scenario.

Finally, the same inflation surge scenario can also be evaluated under stronger anticipation assumptions. In the simulations represented by the dotted lines in Figure 8, we input the same sequence of shocks as previously constructed into the model but now assume that the economy fully anticipates the effects as early as Q1:2020. This illustrates a highly forward-looking extreme case where the pricing pressures are anticipated immediately following the pandemic's opening phase. Under backward expectations, there is no significant economic difference between anticipation and no anticipation. However, under Full MCE, the price shock sequences result in a more gradual inflation increase, with price growth peaking below the initial scenario's peak for Full MCE. Households' decisions are then influenced more immediately, as they assess the post-pandemic inflation surge's implications for their real income early on. The central bank would also begin to tighten policy gradually, in foresight of price pressures building up in the economy.

In summary, these scenario exercises confirm that model-consistent expectations can significantly alter the cyclical properties viewed through the lens of the model, with critical implications for policy assessment. The cases presented are exemplary for some alternative economic repercussions that the central bank may distill when

assessing the implications of supply shocks. To make recommendations for policy, timing, anticipation, and speed of adjustments to supply shocks all matter. Our experiments are also consistent with findings in the literature that, if the monetary authority had had full knowledge of the economic conditions prevailing from 2021 onward, optimal strategies might have called for anticipated hiking of interest rates in response to the increase in inflation (see Darracq Pariès, Kornprobst, and Priftis 2024).

5. Conclusion

ECB-(RE)BASE has been developed as part of the semi-structural model family around ECB-BASE/MC. It is a variant of the euro-area model ECB-BASE that can be simulated under a model-consistent expectation setting. Our implementation allows for the integration of different expectation settings in different parts of the model economy. A key technical challenge of ECB-(RE)BASE relates to the treatment of adjustment costs, which determine the sluggishness of macroeconomic transmissions.

We have shown that the expectation setting can significantly influence the transmission of monetary policy shocks through the model economy. The MCE setting leads to a dampening of monetary policy transmission to the real side of the economy because the modeling of gradual adjustment forces and the prominence of expected lifetime income dominate any implications of the financial sector being forward-looking. The MCE setting also leads to a front-loading of HICP responses due to the exchange rate effect on imported energy. Policy announcements lead to partial anticipation of economic dynamics if credibly understood by the private sector.

Additionally, we have illustrated the implications of expectation settings for macroeconomic volatility through an applied experiment. The markup and negative productivity shocks that have driven up inflation in recent years result in stronger nominal amplifications when assessed through a model with model-consistent expectations compared to a backward-looking setting. Anticipation of shocks increases the complexity of the inflation-output trade-off because the speed and persistence of transmission to the real and to the nominal side may differ substantially.

This information is crucial for the conduct of monetary policy. If the private sector is highly forward-looking, the same sequence of shocks may trigger more pronounced inflation cycles. Consequently, monetary policy needs to respond proactively to effectively limit price pressures while minimizing negative impacts on output.

References

- Adjemian, S., N. Bokan, M. Darracq Pariès, G. Müller, and S. Zimic. 2024. "ECB-(RE)BASE: Heterogeneity in Expectation Formation and Macroeconomic Dynamics." Working Paper No. 2965, European Central Bank.
- Adjemian, S., M. Juillard, F. Karamé, W. Mutschler, J. Pfeifer, M. Ratto, N. Rion, and S. Villemot. 2024. "Dynare: Reference Manual, Version 6." Dynare Working Papers No. 80, CEPREMAP.
- Adolfson, M., S. Laséen, J. Lindé, and M. Villani. 2008. "Evaluating an Estimated New Keynesian Small Open Economy Model." *Journal of Economic Dynamics and Control* 32 (8): 2690–721.
- Angelini, E., N. Bokan, K. Christoffel, M. Ciccarelli, and S. Zimic. 2019. "Introducing ECB BASE: The Blueprint of the New ECB Semi-Structural Model for the Euro Area." Working Paper No. 2315, European Central Bank.
- Angelini, E., M. Damjanović, M. Darracq Pariès, and S. Zimic. 2020. "ECB-BASIR: A Primer on the Macroeconomic Implications of the COVID-19 Pandemic." Working Paper Series No. 2431, European Central Bank.
- Barsky, R. B., and E. R. Sims. 2011. "News Shocks and Business Cycles." *Journal of Monetary Economics* 58 (3): 273–89.
- Beaudry, P., T. J. Carter, and A. Lahiri. 2023. "The Central Bank's Dilemma: Look Through Supply Shocks or Control Inflation Expectations?" Technical Report, National Bureau of Economic Research.
- Bernanke, B. S. 2024. "Forecasting for Monetary Policy Making and Communication: A Review." Technical Report, Bank of England.
- Bernanke, B. S., M. T. Kiley, and J. M. Roberts. 2019. "Monetary Policy Strategies for a Low-Rate Environment." *AEA Papers and Proceedings* 109 (May): 421–26.
- Brayton, F., M. Davis, and P. Tulip. 2000. "Polynomial Adjustment Costs in the FRB-US." Mimeo, Federal Reserve Bank.

- Brayton, F., T. Laubach, and D. L. Reifschneider. 2014. "The FRB/US Model: A Tool for Macroeconomic Policy Analysis." FEDS Notes, Board of Governors of the Federal Reserve System, April 3.
- Brayton, F., E. Mauskop, D. L. Reifschneider, and P. Tinsley. 1997. "The Role of Expectations in the FRB/US Macroeconomic Model." *Federal Reserve Bulletin* 83 (April): 227–45.
- Brayton, F., and P. Tinsley. 1996. "A Guide to FRB/US: A Macroeconomic Model of the United States." Finance and Economics Discussion Series No. 96-42, Board of Governors of the Federal Reserve System.
- Christoffel, K., F. Mazelis, C. Montes-Galdón, and T. Müller. 2022. "Disciplining Expectations and the Forward Guidance Puzzle." *Journal of Economic Dynamics and Control* 137 (C): 104336.
- Chung, H., J.-P. Laforte, D. L. Reifschneider, and J. C. Williams. 2012. "Have We Underestimated the Likelihood and Severity of Zero Lower Bound Events?" *Journal of Money, Credit and Banking* 44 (s1): 47–82.
- Cochrane, J. H. 2017. "The New-Keynesian Liquidity Trap." *Journal of Monetary Economics* 92: 47–63.
- Darracq Pariès, M., A. Kornprobst, and R. Priftis. 2024. "Monetary Policy Strategies to Navigate Post-Pandemic Inflation: An Assessment Using the ECB's New Area-Wide Model." Working Paper No. 2935, European Central Bank.
- Fuller, R., and C.-C. Hsia. 1984. "A Simplified Common Stock Valuation Model." *Financial Analysts Journal* 40 (5), September: 49–56.
- Kennan, J. 1979. "The Estimation of Partial Adjustment Models with Rational Expectations." *Econometrica* 47 (6): 1441–55.
- Kiley, M., and J. Roberts. 2017. "Monetary Policy in a Low Interest Rate World." *Brookings Papers on Economic Activity* 48 (Spring): 317–72.
- Laubach, T., and D. L. Reifschneider. 2003. "Household Consumption and Investment in the FRB/US Model." Mimeo, Board of Governors of the Federal Reserve System.
- Maliar, L., and J. B. Taylor. 2024. "Odyssean Forward Guidance in Normal Times." *Journal of Economic Dynamics and Control* 165 (August): 104877.

- Reifschneider, D. L., and J. M. Roberts. 2006. "Expectations Formation and the Effectiveness of Strategies for Limiting the Consequences of the Zero Bound." *Journal of the Japanese and International Economies* 20 (3): 314–37.
- Reifschneider, D. L., and J. C. Williams. 2000. "Three Lessons for Monetary Policy in a Low-Inflation Era." *Journal of Money, Credit and Banking* 32 (4): 936–66.
- Rotemberg, J. J. 1982. "Monopolistic Price Adjustment and Aggregate Output." *Review of Economic Studies* 49 (4): 517–31.
- Schmitt-Grohé, S., and M. Uribe. 2012. "What's News in Business Cycles." *Econometrica* 80 (6): 2733–64.
- Slobodyan, S., and R. Wouters. 2012. "Learning in an Estimated Medium-Scale DSGE Model." *Journal of Economic Dynamics and Control* 36 (1): 26–46.
- Tetlow, R. J. 2022. "How Large is the Output Cost of Disinflation?" Finance and Economics Discussion Series No. 2022-079, Board of Governors of the Federal Reserve System.
- Tinsley, P. A. 2002. "Rational Error Correction." *Computational Economics* 19: 197–225.
- Warne, A. 2023. "DSGE Model Forecasting: Rational Expectations vs. Adaptive Learning." Working Paper No. 2768, European Central Bank.
- Woodford, M. 2001. "Fiscal Requirements for Price Stability." *Journal of Money, Credit and Banking* 33 (3): 669–728.

QE: Implications for Bank Risk-Taking, Profitability, and Systemic Risk*

Supriya Kapoor^a and Adnan Velic^b

^aTrinity College Dublin

^bTechnological University Dublin

In the aftermath of the subprime mortgage bubble, the Federal Reserve implemented large-scale asset purchase (LSAP) programs that aimed to increase bank liquidity and lending. The excess liquidity created by quantitative easing (QE), in turn, may have stimulated bank risk-taking in search of higher profits. Using comprehensive data on balance sheets, risk measures, and daily market returns in the U.S., we investigate the link between QE, bank risk-taking, profitability, and systemic risk. We find heterogeneous effects across different rounds of QE. In particular, during the third round of QE, banks that were more exposed to unconventional monetary policy through mortgage-backed securities (MBS) purchases increased their risk-taking behavior and profitability. However, these banks also reduced their contribution to systemic risk, indicating that the implementation of QE had an overall stabilizing effect on the banking sector. These results highlight the different distributional effects of QE.

JEL Codes: E52, E58, G21.

1. Introduction

How does quantitative easing (QE) affect individual banks and overall financial stability? Our paper seeks to examine various dimensions of this question. Central banks have traditionally employed

*We thank seminar and conference participants at ECB Research Workshop, De Nederlandsche Bank, Sveriges Riksbank, UC Santa Cruz seminar series, Irish Economic Association, Baltic Economic Association, and FEBS for helpful comments. We are grateful to the editor, Christopher Waller, and an anonymous referee for their helpful comments. Author e-mails: supriya.kapoor@tcd.ie and adnan.velic@tudublin.ie.

short-term interest rates to stabilize inflation and economic fluctuations. The period preceding the Global Financial Crisis (GFC) of 2007–08 featured low monetary policy rates that led to a deterioration of bank lending standards and subsequent increases in risk-taking (Dell’Ariccia, Laeven, and Marquez 2014; Bernanke et al. 2020). The persistently low real interest rate environment fueled a boom in asset prices and securitized credit that led financial institutions to take on increasing risk and leverage, thereby extending loans to borrowers with minimum or no credit history (Altunbas, Gambacorta, and Marques-Ibanez 2010; Paligorova and Santos 2017). Hence, banks are considered to have played a significant role in the buildup and spread of systemic risk (Nijskens and Wagner 2011). The Financial Stability Board defines systemic risk as “a risk of disruption to financial services that is caused by an impairment of all or parts of the financial system and has the potential to have serious negative consequences for the real economy” (Financial Stability Board 2009).

The Federal Reserve (Fed) responded to the zero lower bound and subsequent worsening of the economic crisis by embarking on large-scale asset purchase programs, often referred to as QE, with the intention of providing additional monetary stimulus, restoring stability, and reducing systemic risk. The introduction of such unconventional policy measures rekindled debate on the impact of monetary policy on risk-taking, performance, and stability of the banking sector. Empirical evidence suggests that QE can have both benefits and costs with varying effects on economic outcomes.¹ On the one hand, QE can induce banks to increase their appetite for risk by extending loans to borrowers with bad credit history or no credit history at all (i.e., risky borrowers). Alternatively, QE may reduce bank risk by increasing bank lending that boosts investment and economic growth. This strengthens banks’ resistance to exogenous shocks. Thus, the net effect of QE on banking-sector stability is not

¹While quantitative easing has shown to boost lending and liquidity in the banking sector (Rodnyansky and Darmouni 2017; Chakraborty, Goldstein, and Mackinlay 2020; Luck and Zimmermann 2020), this can also come at the cost of increased risk-taking by banks (Kandrac and Schlusche 2021; Kurtzman, Luck, and Zimmermann 2022).

obvious and depends on whether the benefits of QE outweigh its costs.

The key question that we ask is: What is the impact of QE on individual bank risk-taking, bank profitability, and systemic risk? Although recent empirical studies examine the effects of monetary policy on individual bank risk,² we are aware of relatively few that address systemic risk. So far the literature is split on the role of bailouts in either increasing or decreasing systemic risk. Evidence suggests that bailouts can reduce systemic risk either through a higher charter value effect³ (Cordella and Yeyati 2003) or through a reduction in undiversifiable contagion risk across banks (Diamond and Rajan 2005; Dell’Ariccia and Ratnovski 2013; Choi 2014). Others predict that bailouts can increase systemic risk by exacerbating moral hazard problems and creating strategic complementarities among banks, which encourage coordinated risk-taking behavior (Acharya and Yorulmazer 2007a, 2007b; Diamond and Rajan 2009a; Farhi and Tirole 2012). Berger, Roman, and Sedunov (2017) employ the U.S. Troubled Assets Relief Program (TARP) as an event study and provide discussion on the topic. To this end, our study is one of the first to provide a distributional perspective on the implications of QE and its effect on systemic risk. Our analysis moreover explores potential asymmetries in the effects of different rounds of QE by the Fed. We focus in particular on purchases of mortgage-backed securities (MBS) since this was the primary asset tool employed. The idea is that such purchases should directly improve the risk profile and liquidity of banks’ overall assets. According to the Federal Open Market Committee (FOMC), based on FRB/US model simulations, purchases of agency MBS were more likely to provide greater economic stimulus than purchases of long-term Treasuries of comparable size (Gagnon and Holscher 2008). The former, therefore, were expected to be of central importance.⁴

²See Altunbas, Gambacorta, and Marques-Ibanez (2010), Delis and Kouretas (2011), Dell’Ariccia, Laeven, and Marquez (2014), and Bikker and Vervliet (2018).

³A bank’s charter value is an indicator of the bank’s incentive to take risk or private cost of failure. See Keeley (1990).

⁴In unreported additional checks, we find that purchases of Treasuries by the Fed did not have the same strong impact on risks and profitability in the banking sector as MBS purchases. Indeed, Krishnamurthy and Vissing-Jorgensen (2011)

We study the heterogeneous impact of QE on bank risk-taking, profitability, and systemic risk using a sample of U.S. bank holding companies over the 2006:Q1–2014:Q4 period. We first assess the impact of QE on bank risk-taking by employing time-varying Z-scores, a classic balance sheet measure of firm riskiness. Second, we explore whether banks increase their risk appetite to lend more during various QE episodes in search of higher profits. Finally, we examine the effects of QE on systemic risk within the U.S. financial sector. Section 3 discusses the three variables of interest in detail.

To identify the effects of QE, we closely follow the empirical strategy of Rodnyansky and Darmouni (2017) and Luck and Zimmermann (2020). We exploit the cross-sectional variation across holdings of MBS for commercial banks in our sample. The identification rests on the idea that banks with higher MBS holdings benefit more from the Fed's asset purchases of securities.⁵ We adopt a fixed-effects identification strategy that relies on the interaction of cross-sectional variation among banks in their MBS holdings and the corresponding QE time dummies. We employ several definitions based on the ratio of MBS-to-total-assets prior to QE in order to classify banks into treated and control groups, respectively, and investigate the differential effects of the policy across banks.

Our baseline statistics offer preliminary insight into the interrelations of the three variables of interest. We then run separate regressions for risk-taking, profitability, and systemic risk, and find that banks with higher MBS holdings increased risk-taking during the third round of QE (QE3) relative to banks that were unaffected by large-scale asset purchases.⁶ This is consistent with the previous literature where Gambacorta (2009), Altunbas, Gambacorta,

conclude that focusing primarily on Treasuries as a policy target is inappropriate, as QE functions through various channels that affect specific asset classes differently.

⁵There are several reasons why banks that held more mortgage-backed securities benefited more from the large-scale asset programs. First, during the three waves of QE, the Fed focused on easing the deterioration in the MBS market by lowering yields and increasing the prices of banks' current asset holdings, thereby improving the balance sheets of banks that held higher shares of mortgage-backed securities. Second, banks with more MBS sold to the Fed saw a higher increase in reserves, which should have shifted their loan supply (Kandrac and Schlusche 2021).

⁶We note that risk-taking is not necessarily undesirable, unlike risk-shifting.

and Marques-Ibanez (2010), Delis and Kouretas (2011), and more recently Kandrac and Schlusche (2021) find a negative relation between monetary policy (conventional and unconventional) and bank risk-taking, thus confirming the existence of the bank risk-taking channel.

Furthermore, we show that the same banks enjoyed higher profits during QE3 and suggest a possible connection between increased bank risk-taking and profit maximization. Finally, we document the effects of QE on systemic risk and financial stability. Our evidence suggests that banks with relatively more MBS to total assets reduce their contributions to systemic risk after the implementation of QE3. This implies that asset purchase programs were supportive in the recovery of the banking sector after the GFC. Our results are consistent with Berger, Roman, and Sedunov (2017), who investigate the impact of bailouts on banks' contribution to overall financial stability. With our findings, we suggest a mechanism of monetary policy transmission where banks reduce their contribution to systemic risk as a result of high risk-taking and subsequent increases in profits. Our empirical findings are robust to alternative methodologies, including a system of equations approach, and other robustness tests.

The remainder of the paper is organized as follows. Section 2 describes our theoretical framework, followed by Section 3, which discusses the data and identification strategy. Section 4 presents our results, while Section 5 concludes.

2. Literature Overview

In this section, we review the theoretical implications of monetary expansion for (i) bank risk-taking, (ii) bank profitability, and (iii) systemic risk.

2.1 *Bank Risk-Taking*

The effect of monetary policy on banks' risk-taking behavior is of special interest, as excessive risk-taking by commercial banks is considered to be one of the factors that led to the outbreak of the

GFC.⁷ Borio and Zhu (2012) refer to the “*risk-taking channel*” of monetary policy as prolonged periods of low interest rates that can alter the risk perceptions and risk tolerance of financial institutions. The channel operates on the basis of valuations, incomes, and cash flows, where an expansionary monetary policy drives up the value of real and financial collateral and increases bank capacity to extend risky loans.

The empirical literature on the risk-taking channel has been widely studied by examining key policy rates. Several studies, including those of Dell’Ariccia and Ratnovski (2013), Ioannidou, Ongena, and Peydró (2014), Jiménez et al. (2014), and Paligorova and Santos (2017), show that banks loosen lending standards and increase credit supply by extending loans to risky borrowers during monetary expansions. The literature adopts several measures for risk-taking ranging from internal ratings on borrowers’ repayment capacity, default probabilities, time to default (Ioannidou, Ongena, and Peydró 2014), and change in expected default frequency (Gambacorta 2009) to Z-scores (Laeven and Levine 2009; Schaeck and Čihák 2010; Beck, De Jonghe, and Schepens 2013). Recent work on the topic is concerned with the impact of LSAPs on the risk-taking behavior of commercial banks. In their empirical analyses, Kandrac and Schlusche (2021) and Kurtzman, Luck, and Zimmermann (2022) use microdata (bank and loan level) and find that various QE programs led to higher total loan growth and an increase in the share of riskier loans within banks’ portfolios.

There are several reasons why QE may have promoted risk-taking. As a response to QE and excess reserves in balance sheets, banks may choose to maintain their profit margins and incomes by searching for yield in lending (Borio and Zhu 2012; Jiménez et al. 2014; Kandrac and Schlusche 2021). Since the main motive of QE was to directly lower long-term interest rates and thus stimulate the economy, lower yields incentivize banks to accelerate risk-taking activity (Rajan 2006). The various QE policies might have induced

⁷Studies such as those of Acharya and Richardson (2009) and Diamond and Rajan (2009b) assert that very long periods of low interest rates, accompanied by an abundance of global liquidity, would have led financial companies to behave in a riskier way. This was indeed a contributing factor in the development of the Global Financial Crisis.

financial market participants to reallocate portfolios toward riskier and more profitable projects, thereby leading to excessive risk-taking (Abbate and Thaler 2019). Finally, monetary policy can affect bank risk-taking through leverage decisions. Dell’Ariccia, Laeven, and Marquez (2014) present a theoretical model in which a reduction in the risk-free interest rate reduces the cost of holding required reserves, thus leading to higher leverage and increased risk-taking by banks. However, directly after the GFC, it is conceivable that banks may have used the initial stages of QE (e.g., first round) to improve or stabilize their own balance sheet risk profiles by not engaging in risky lending/investment, with a view to restoring confidence in their operations. Asymmetric effects on bank risk-taking across different QE rounds may also be expected if yields are differentially affected across these phases. Krishnamurthy and Vissing-Jorgensen (2011, 2014), for instance, demonstrate that the second round of QE had very weak effects on yields compared to those of the third round. The authors also illustrate that MBS yields were more strongly affected than Treasury yields by QE, suggesting that Fed purchases of MBS had more potent effects on banks than Treasury purchases. They contend that purchases of the former are crucial in attenuating corporate credit risk.

As the Fed buys assets from banks through QE, it increases assets (securities) on its balance sheet, concurrently crediting banks’ reserve accounts. By providing liquidity in the banking sector, QE made it easier and cheaper for banks with relatively high MBS to extend loans to nonfinancial firms and households, thereby inducing banks to increase their risk-taking capacity.

2.2 Bank Profitability

While risk-taking behavior is an intended consequence of unconventional monetary policy, reach for yield is a form of risk-taking, a way to achieve higher profits and performance. An expansionary monetary policy induces banks to increase their holdings of risky assets in order to “reach for yield” (Rajan 2006; Borio and Zhu 2012). Reaching for yield means that banks maximize the yield spread without paying adequate regard to risk—by obtaining a higher return on their asset holdings, by paying a lower return to their creditors, or by engaging in both practices. In other words, it can be defined as

the agents' propensity to buy riskier assets in order to achieve higher yields and profits.⁸

Interest in the relation between monetary policy and bank profitability has gained momentum since the GFC. Accommodative monetary policy creates bank liquidity and lowers the cost of debt, thereby increasing bank capital, lowering loan loss provisioning, and increasing profitability (Bernanke and Gertler 1995; Bernanke 2007; Freixas, Martin, and Skeie 2011). At the same time, periods of monetary expansion can have negative effects on bank profits. Additional monetary stimulus when interest rates are already too low can squeeze banks' net interest (profit) margins, in turn diminishing their net worth and credit supply. Persistently low interest rates may place chronic downward pressure on profitability that eventually more than offsets banks' initial capital gains and likewise hinders the credit supply. Importantly, there exists a "reversal interest rate" at which accommodative monetary policy can become contractionary for lending (Abadi, Brunnermeier, and Koby 2023). Inertia at low interest rate levels can thus hamper the monetary policy transmission mechanism (Alessandri and Nelson 2015; Brunnermeier and Koby 2016; Borio, Gambacorta, and Hofmann 2017). In order to improve net interest margins when interest rates are too low for too long, banks might reduce their funding rates to compensate for low lending rates. The net effect of monetary policy measures on bank profitability, therefore, is not obvious and remains an empirical question.

This paper further investigates the impact of QE on bank profitability and, hence, the soundness of the banking sector. Keeley (1990) shows that risk-taking incentives should be lower in more profitable firms, as they stand to lose more shareholder value if downside risks are realized. However, more recently, significant risk-taking in profitable institutions seems to contradict the traditional predictions of corporate finance models (Martynova 2015). In an unconventional monetary policy environment, banks may increase their risk-taking and influence asset valuations in search for higher profits.

⁸Exploring the drivers of relative finance wages, Velic (2025) finds that financial deregulation also contributed to higher profitability and relative skilled wages in the financial sector, especially in the lead-up to the GFC.

2.3 *Systemic Risk*

The literature on unconventional monetary policy thus far focuses on the effects of QE on interest rates (Krishnamurthy and Vissing-Jorgensen 2011), bank lending (Rodnyansky and Darmouni 2017; Luck and Zimmermann 2020), risk-taking (Kandrac and Schlusche 2021), and bank liquidity (Kapoor and Peia 2021), with little attempt to investigate the impact of QE on systemic risk. This paper explores the distributional effects of QE by investigating the impact on systemic risk. In particular, we study if QE led financial institutions to increase their risk-taking and profitability, thus providing a cushion to absorb losses and reduce the likelihood of financial distress and bank failure.

Prior to the crisis, Basel regulations and capital ratios were popular tools to measure a bank's risk. However, since the GFC, use of market-based measures of systemic risk has become more prevalent. For instance, Huang, Zhou, and Zhu (2012) uncover the concept of "distress insurance premium" using credit default swaps and equity returns. Similarly, Adrian and Brunnermeier (2011) construct the "CoVaR" that captures the contribution of an institution to the overall systemic risk of the system. Allen, Bali, and Tang (2012) extend the concept of value at risk (VaR) by constructing CATFIN, which links systemic risk to VaR of the financial system. Acharya, Engle, and Richardson (2012) propose the marginal expected shortfall (MES) that depends solely on market-based information and estimates individual bank reactions to the entire stock market when aggregate returns are low. Furthermore, Acharya et al. (2017) estimate systemic expected shortfall (SES), which merges both market and balance sheet information to measure a bank's propensity to be undercapitalized under stress conditions. We employ SES for our analysis. Details on the construction of the measure are provided in the next section.

3. Data and Methodology

Our data come from the quarterly Consolidated Financial Statements for Bank Holding Companies (BHCs) in the United States, which are available from the Federal Reserve Bank of Chicago. All BHCs are subject to regulation by the Federal Reserve Board of

Governors under the Bank Holding Company Act of 1956 and Regulation Y. Our data cover the period 2006:Q1 to 2014:Q4 and include information on the financial condition of the BHCs such as balance sheet items, off-balance-sheet exposures, statistics on different types of loans, etc.⁹ We omit observations with missing values for total assets. This leaves us with a final sample of 1,438 unique BHCs and 31,467 BHC-quarter observations.

We match these data with market data obtained from the Center for Research in Security Prices (CRSP), which provides U.S. stock data for all publicly listed institutions (including financial firms) with end-of-day and month information on all primary listings (NYSE, NYSE MKT, NASDAQ) along with basic market indices. The data include multiple company identifiers—including PERMNO, PERMCO, and GVKEY—that are used to match with bank-level data. Each BHC's identifier was matched with the CRSP permanent company identifier (PERMCO) using the CRSP-FRB Link provided by the Federal Reserve Bank of New York.¹⁰ In order to fully utilize the two data sets, we employ fuzzy matching for all companies that did not have a PERMCO available.¹¹ Our merged data comprise information on daily returns for 743 unique bank holding companies, with 11,738 BHC-quarter observations in the sample.¹²

To estimate the impact of QE on risk-taking and systemic risk, several key measures are constructed separately from estimated

⁹The reporting forms have changed a number of times over the sample period, causing changes to some variables available in the raw data over time. Where reporting changes have affected variables of interest, we have manually traced these changes through the reporting form and merged data as appropriate.

¹⁰According to the Federal Reserve Bank of New York, the CRSP and BHC data, when merged using PERMCOs, will yield 1,430 unique BHCs from June 30, 1986 to December 31, 2017. Available at https://www.newyorkfed.org/research/banking_research/datasets.html.

¹¹These companies were matched based on their names, with a threshold level of 0.95. The matching was further cross-checked by employing different threshold levels (between 0.95 and 1) in order to ascertain the quality of the matching method.

¹²The matched data are used in analyzing the impact of QE on systemic risk contribution by financial firms. Thus, we end up with 288 BHCs and approximately 5,100 BHC-quarter observations after the construction of systemic risk measures.

regressions. First, we compute Z-scores as a measure of bank risk-taking by following Schaeck and Čihák (2010) and Beck, De Jonghe, and Schepens (2013). The Z-score is a classic balance sheet measure of bank riskiness, conceived from the concept of a bank's probability of default. It is often used to capture either the stability of the banking sector (Lee and Hsieh 2014) or the inverse probability of insolvency of a bank (Williams 2016). The Z-score is formally expressed as

$$Z_{i,t} = \frac{ROA_{i,t} + EA_{i,t}}{\sigma_{i,t}^{ROA}}, \quad (1)$$


where $ROA_{i,t}$ is the return on assets for bank i at time t (bank profitability), $EA_{i,t}$ is the ratio of bank i 's equity to total assets at time t (capitalization), and $\sigma_{i,t}^{ROA}$ is the time-varying variability of returns on assets. A lower Z-score indicates higher bank risk-taking and vice versa. As Z-score data are highly skewed, we employ the natural logarithm of the measure. For brevity, we refer to the natural logarithm of the Z-score as "Z-score" throughout the study. Figure 1 plots the components of the average Z-score over time. Based on gross correlations, the Z-score is most strongly related to the ratio of equity to assets. Second, we examine net interest income and return on assets as measures of profitability, with the latter measure featuring in our robustness checks.

Third, we employ the systemic expected shortfall (SES) measure proposed by Acharya et al. (2017) for systemic risk. SES captures the contribution of individual banks to systemic risk in the banking sector. The measure employs both market and balance sheet information and gauges a bank's propensity to be undercapitalized under stress conditions. Therefore, marginal expected shortfall (MES) and a financial firm's leverage are two significant components in the construction of SES.¹³ We calculate MES by taking average returns for a firm on days when the market as a whole is in the tail of its loss distribution:

$$MES_{i,t} = E[R_t^i | R_t^m < C]. \quad (2)$$

¹³MES depends solely on market-based information and estimates an individual bank's reaction to the entire stock market when aggregate returns are low.

Figure 1. Components of Z-score



Next, leverage is defined as the quasi-market value of assets to market value of equity and is estimated by the standard approximation approach using book liabilities and market equity:

$$LVG_{i,t} = \left[\frac{(BookAssets_{i,t} - BookEquity_{i,t}) + MarketEquity_{i,t}}{MarketEquity_{i,t}} \right], \tag{3}$$

where $MarketEquity_{i,t}$ is the number of shares outstanding times the average price of a share for a given quarter, and $BookAssets_{i,t}$ and $BookEquity_{i,t}$ are bank characteristics that are available at a quarterly frequency from balance sheet data. Using $MES_{i,t}$ and $LVG_{i,t}$, Acharya et al. (2017) regress the percentage stock returns of U.S. institutions on the given components prior to the crisis and employ an $SES_{i,t}$ definition of the following form:

$$SES_{i,t} = \hat{a}MES_{i,t} + \hat{b}LVG_{i,t}. \tag{4}$$

We adopt a similar formula and take values of coefficient estimates \hat{a} and \hat{b} from Acharya et al. (2017).¹⁴ $SES_{i,t}$ measures the extent to which a bank is undercapitalized in an event in which the entire financial system is under distress. Thus, increases in $SES_{i,t}$ indicate increases in banks' expected losses during crisis.

In order to identify the effects of QE on bank risk-taking, profitability, and systemic risk, we employ the fixed-effects estimation methodology and divide our sample into treatment and control groups based on the ratio of MBS-to-total-assets prior to the introduction of QE, i.e., 2007:Q4.¹⁵ Our identification strategy follows Rodnyansky and Darmouni (2017) and exploits cross-sectional variation in MBS holdings across banks. The identification strategy relies on the assumption that banks with higher proportions of MBS to total assets were more affected by the asset purchases. This is corroborated by the fact that banks with more MBS sold to the Fed saw a higher increase in reserves, which should have shifted their loan supplies (Kandrac and Schlusche 2021).


We measure a bank's exposure to QE by the ratio of MBS to total assets.¹⁶ We define treatment group banks as those in the top 25 percent (top quartile) of the MBS-to-total-assets distribution, while control group banks are in the bottom 25 percent (bottom quartile) of the distribution. Figure 2 shows the evolution of the MBS to total assets of treated and control banks separately, indicating that MBS holdings of treated banks (solid line) start to decline immediately after the implementation of QE, while control banks (dashed line) see an increase. We use several other definitions of treatment and control groups based on deciles (top and bottom 10 percent) and the continuous ratio (full sample) of MBS to total assets in 2007:Q4.

¹⁴ Additionally, we examine our own estimates of coefficients a and b in unreported robustness checks. The results are similar qualitatively. MES and LVG tend to exhibit an inverse contemporaneous relation with SES .

¹⁵ The choice of time period 2007:Q4 minimizes endogeneity since this is more than six months before QE1. The MBS threshold is defined in a particular quarter so that we are able to track the effects for the same set of banks over time. Implementing the treatment definition in alternative quarters just before each of the other two QE episodes yields similar results qualitatively.

¹⁶ We also employ the MBS-to-securities ratio as an alternative measure of a bank's exposure to QE. The results are consistent with the main results since the ratios of MBS to assets and MBS to securities are highly correlated in the sample.

Figure 2. MBS to Total Assets for Treated and Control Banks



Note: The figure maps the evolution of the ratio of MBS to assets for treated and control banks. Treated banks are banks in the top quartile (top 25 percent) of the MBS-to-total-assets ratio in 2007:Q4, while control banks are in the bottom quartile (bottom 25 percent) of the distribution. Shaded areas highlight the three episodes of QE.

We employ a fixed-effects identification strategy that relies on the interaction of cross-sectional variation among banks in their MBS holdings and the corresponding QE time dummies. Our regression model is given by

$$\begin{aligned}
 Y_{i,t} = & \alpha_i + \beta_t + \rho_{i,t} + \gamma'_1 QE_t + \gamma_2 Treat_i + \theta' Treat_i \\
 & \times QE_t + \delta' X_{i,t-1} + \epsilon_{i,t}.
 \end{aligned}
 \tag{5}$$

$Y_{i,t}$ denotes the dependent variable and is the measure of individual bank risk-taking, profitability, or contribution to systemic risk across our three specifications. $Treat_i$ is an indicator variable and takes the value of 1 if a bank belongs to the treatment group and 0 otherwise. $QE'_t = [QE1_t, QE2_t, QE3_t]$ is a vector of time dummies corresponding to the introduction of each QE episode, where

the dummy takes the value of 1 during the episode and 0 otherwise. QE1 corresponds to the period 2008:Q4 (November)–2010:Q2 (June), QE2 to 2010:Q4 (November)–2011:Q2 (June), and QE3 to 2012:Q3 (September)–2014:Q3 (October). $Treat_i \times QE_t$ is an interaction term between a bank's treatment status and time dummies corresponding to each QE episode. The vector θ captures our coefficients of interest, namely the differential impact of each round of QE on the dependent variable in the treated as compared to the control group.¹⁷

Vector $X_{i,t-1}$ includes a series of bank-level controls that capture differences in the scale and financial position of banks that might affect their activities. Specifically, we control for bank size (natural logarithm of total assets), liquidity ratio (ratio of cash to total deposits), tier 1 capital ratio to account for the extent to which a bank absorbs potential losses, leverage ratio, and deposit ratio. All control variables are included as lagged terms, and a description of the construction of the same is presented in Table A.1 in Appendix A. Including interactions between controls and the treatment variables does not alter results significantly.

The regression also includes bank fixed effects (α_i) to control for fixed differences among banks, and quarterly time fixed effects (β_t) to control for residual intertemporal differences in common shocks that affect all banks.¹⁸ Examples of shocks affecting all banks during the sample period are the Dodd-Frank Act and Basel Accords (including stress-testing). In addition, we employ state fixed effects ($\rho_{i,t}$) to control for location fixed effects for BHCs as well as any state-specific regulatory policies that might affect banking functions.

While the null hypothesis of cross-sectional independence can be rejected in tests for all regression model residuals, we report in table notes that average and average absolute cross-sectional correlations are quite small, thus making the assumption of independence appropriate in practice (Galstyan and Velic 2017; Pesaran 2021). The summary statistics for our data set and the key variables employed in the study can be seen in Table 1.

¹⁷The individual effects of the treatment variable are absorbed by fixed effects in regressions.

¹⁸Statistically significant time dummies are included in the regression, and there are no multicollinearity issues in the specification.

Table 1. Summary Statistics

Variable	Mean	Standard Deviation	p25	p50	p75	Observations
Treatment Variable: MBS/Total Assets	0.095	0.088	0.026	0.076	0.138	31,754
Dependent Variables: ln(<i>Z-score</i>)	3.38	0.687	3.05	3.417	3.752	27,094
Risk Assets/Assets	0.933	0.064	0.918	0.953	0.972	31,754
ln(Net Interest Income)	10.11	1.37	9.26	9.87	10.61	31,754
Return on Assets	0.095	9.465	0.002	0.005	0.008	28,508
$\Delta \ln(SEI)$	-3.1	1.00	-3.3	-3.25	-3.19	5,087
<i>SRIISK</i>	-3.34	13.21	0.63	0.82	1.01	4,843
Bank-Specific Controls: Bank Size	14.176	1.325	13.365	13.768	14.534	31,754
Tier 1 Capital Ratio	13.932	22.608	10.67	12.57	15.03	30,484
Leverage Ratio	9.968	15.371	8.19	9.31	10.63	30,484
Deposits Ratio	0.782	0.113	0.750	0.805	0.849	29,408
Liquidity	0.854	65.42	0.029	0.045	0.083	29,388

Note: Summary statistics recorded from 2006:Q1 to 2014:Q4 for all U.S. BHCs. All variables are at quarterly frequency. Variable definitions are provided in Appendix A. Risk assets/assets, return on assets, and *SRIISK* are alternative measures of bank risk-taking, profitability, and systemic risk, respectively, in robustness checks.

Table 2. Pearson Gross Correlations

Variable	Z-score	NII	SES
Z-score	1.0000		
NII	-0.0599***	1.0000	
SES	-0.1593***	0.0119	1.0000

Note: The variables Z-score and net interest income (NII) are expressed in natural logarithms, while SES is defined as the change in systemic expected short-fall. ***, **, and * represent significance at 1 percent, 5 percent, and 10 percent, respectively.

4. Results

4.1 Baseline Correlations

We present pairwise (Pearson) gross correlations between bank risk-taking, profitability, and systemic risk in Table 2. Although relatively weak, we report a negative correlation between the Z-score and net interest income, indicating that a lower Z-score (higher bank risk-taking) is associated with higher net interest income (profitability). Similarly, we observe a negative gross relation between the Z-score and systemic risk, while bank profitability and systemic risk display an insignificant positive correlation.

Table 3 presents partial correlations between the three aforementioned dependent variables. The partial correlation between risk-taking and profitability in column 1 is consistent with the corresponding gross correlation in Table 2. However, once we take into account control factors, the coefficient between systemic risk and bank risk-taking turns positive, as shown in column 3 of Table 3, indicating that higher risk-taking (lower Z-score) is associated with lower systemic risk. Finally, the same table reports that higher bank profitability is now associated with lower systemic risk (column 2). The results suggest that gross relations can obscure the picture of interlinkages. The findings in the subsequent sections help us to build on these correlations and establish a distinct channel that highlights the distributional effects of QE on risk-taking, profitability, and systemic risk.

Table 3. Baseline Partial Correlations

	Z-score	SES	
	(1)	(2)	(3)
Net Interest Income	-0.091* (0.050)	-0.052* (0.027)	
Z-score			0.023** (0.009)
Size	0.082 (0.050)	0.055** (0.028)	0.003 (0.006)
Liquidity	-1.554*** (0.091)	0.316*** (0.058)	0.466*** (0.081)
Tier 1 Capital Ratio	0.023*** (0.005)	0.005** (0.003)	0.001 (0.004)
Leverage Ratio	0.065*** (0.007)	-0.018*** (0.004)	-0.022*** (0.006)
Observations	4,461	4,981	4,473
R-squared	0.919	0.204	0.198
Bank Fixed Effects	Yes	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes

Note: The dependent variable in column 1 is the logarithm of Z-score, while in columns 2 and 3 it is the systemic expected shortfall. The null hypothesis of cross-sectional independence in Pesaran's (2021) test can be rejected for regression residuals at conventional significance levels. However, the average and average absolute cross-sectional correlations are relatively low, standing approximately at 0.1 and 0.3 for risk-taking (Z-scores) and approximately 0.04 and 0.2 (columns 2 and 3) for systemic risk (SES). According to Pesaran (2021), in finite samples, the assumption of cross-sectional independence is appropriate when cross-sectional correlations are low. Constant terms included but not reported. Robust standard errors in parentheses. ***, **, and * represent significance at 1 percent, 5 percent, and 10 percent, respectively.

4.2 Benchmark QE Regressions

4.2.1 Bank Risk-Taking

Table 4 presents the main regression results for the impact of QE on individual bank risk-taking. Corresponding to different definitions of the treatment variable, columns 2, 4, and 6 present results using

Table 4. The Impact of QE on Bank Risk-Taking

	Z-score					
	(1)	(2)	(3)	(4)	(5)	(6)
$QE1_t \times Treat_i^Q$	0.010 (0.018)	0.022 (0.017)				
$QE2_t \times Treat_i^Q$	0.008 (0.025)	-0.013 (0.023)				
$QE3_t \times Treat_i^Q$	-0.043** (0.017)	-0.048*** (0.017)				
$QE1_t \times Treat_i^D$			-0.026 (0.027)	-0.011 (0.025)		
$QE2_t \times Treat_i^D$			0.008 (0.038)	-0.040 (0.035)		
$QE3_t \times Treat_i^D$			-0.054** (0.026)	-0.125*** (0.026)		
$QE1_t \times \left(\frac{MBS}{TotalAssets}\right)_i$					0.074 (0.085)	0.104 (0.084)
$QE2_t \times \left(\frac{MBS}{TotalAssets}\right)_i$					0.098 (0.117)	0.050 (0.115)
$QE3_t \times \left(\frac{MBS}{TotalAssets}\right)_i$					-0.220*** (0.081)	-0.263*** (0.087)

(continued)

Table 4. (Continued)

	Z-score					
	(1)	(2)	(3)	(4)	(5)	(6)
Observations	11,391	10,128	4,591	4,082	20,876	19,724
R-squared	0.077	0.114	0.067	0.102	0.122	0.127
Bank-Level Controls	No	Yes	No	Yes	No	Yes
Bank Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Note: The dependent variable in columns 1–6 is the logarithm of Z-scores. $Treat_i^Q$ takes the value 1 for banks in the top quartile (top 25 percent) of the MBS-to-total-assets distribution, while 0 for banks in the bottom quartile (bottom 25 percent) of the distribution. $Treat_i^D$ takes the value 1 for banks in the top decile (top 10 percent) of the MBS-to-total-assets ratio, and 0 for banks in the bottom decile (bottom 10 percent) of the distribution. $\left(\frac{MBS}{TotalAssets}\right)_i$ is the ratio of MBS to total assets in 2007:Q4. $QE1_t$, $QE2_t$, and $QE3_t$ are dummy variables corresponding to the respective QE waves, where the dummy takes the value 1 during the QE episode and 0 otherwise. Control variables include bank size, liquidity, tier 1 capital ratio, loans-to-assets ratio, and leverage ratio. All control variables are taken as lagged. Including interactions between controls and the treatment variables does not alter results significantly. The null hypothesis of cross-sectional independence in Pesaran’s (2021) test can be rejected for regression residuals at conventional significance levels. However, the average and average absolute cross-sectional correlations are relatively low, each standing approximately at 0.1 across regressions. According to Pesaran (2021), in finite samples, the assumption of cross-sectional independence is appropriate when cross-sectional correlations are low. Constant terms included but not reported. Robust standard errors in parentheses. ***, **, and * represent significance at 1 percent, 5 percent, and 10 percent, respectively.

the main specification model that includes control variables. Meanwhile, columns 1, 3, and 5 exclude all bank controls. The regression employs the Z-score as the dependent variable, which indicates the inverse probability of insolvency. A lower Z-score represents higher risk-taking by banks.

The results suggest that banks with larger MBS holdings increase their relative risk-taking capacity during the third round of QE (QE3). This is indicated by the negative and statistically significant coefficient on the interaction term between treated banks and QE3 across most specifications. In particular, banks in the top quartile and decile, respectively, of the MBS-to-total-assets ratio show an increase in risk-taking by 4.8 percent (column 2) and 12.5 percent (column 4), respectively. Additionally, we employ the treatment status of banks based on the continuous measure of MBS to total assets. It can be inferred from the results shown in column 6 that banks with higher MBS to assets increased their risk-taking by approximately 26.3 percent.¹⁹

Our results complement the findings of Luck and Zimmermann (2020) that show higher loan supply by the treated banks during QE3, a possible channel for increased riskiness as shown in Kurtzman, Luck, and Zimmermann (2022).²⁰ Our results are also consistent with the previous literature on the impact of conventional monetary policy on bank risk-taking, including Gambacorta (2009), Altunbas, Gambacorta, and Marques-Ibanez (2010), and Delis and Kouretas (2011) among others, that confirms the presence of a risk-taking channel during episodes of expansionary monetary policy.

Table 4 further indicates statistically insignificant effects for bank risk-taking during QE1. The direction of the coefficients and their economic significance is opposite to the evidence obtained for QE3.²¹ Ignoring statistical significance, the results overall suggest that bank

¹⁹More accurate effects can be obtained from $(e^\theta - 1) \times 100$.

²⁰Although Kurtzman, Luck, and Zimmermann (2022) highlight the negative relation between QE and bank lending standards, which measures the quality of loans rather than quantity, their findings address the fact that banks increase their risk-taking due to both lower lending standards and greater availability of credit.

²¹Since Z-score is a measure of bank stability, note that a positive coefficient on the interaction term would imply high stability and low risk-taking by treated banks during QE1.

risk-taking decreased in the initial phases of QE. Taking statistical significance into account, the interpretation is that banks refrained from additional risk-taking despite an increase in liquidity, implying a decrease in relative risk-taking when compared to the amount of liquidity available. Our understanding of this result is based on the study of Kapoor and Peia (2021), which yields no evidence of loan growth during QE1. Since studies such as that of Kurtzman, Luck, and Zimmermann (2022) associate riskiness with bank lending, an increase in riskiness during QE1 was an unlikely result. Moreover, the QE3 program was larger in scale as compared to QE1, implying that banks, as a response to the economic crisis of 2007–08, used the QE program in the first round to reduce their riskiness.²² By the time QE3 was implemented, banks were better prepared to increase their appetite for risk through higher lending. Our results reveal that it was only after QE3, as opposed to QE1, that banks' portfolios reflected this pattern. This is consistent with Luck and Zimmermann (2020), who investigate the employment effects of disparate QE rounds and find a substantial difference between treated and control banks after QE3 only.

The regression controls for time-variant bank characteristics include bank size, tier 1 capital ratio, liquidity, leverage ratio, and loans-to-asset ratio. The overall findings are relatively robust across all specifications. We note that the time-invariant heterogeneity between treated and control groups is addressed via bank fixed effects, while the demand-side conditions that affect borrower riskiness are controlled by employing year-quarter and state fixed effects.

4.2.2 Bank Profitability

The results in Table 5 investigate the impact of QE on bank profitability, measured by the logarithm of net interest income which reflects the capability of a bank to generate profits from its asset management functions.²³ As before, the main variables of interest

²²This is somewhat implied by the statistically insignificant positive coefficients on the interaction terms in columns 1–2 and 5–6 for QE1.

²³We lose only nine observations by taking the logarithm of net interest income, which is a minuscule amount given that we have over 30,000 observations in the sample. Using the non-log variable yields similar results qualitatively.

Table 5. The Impact of QE on Bank Profitability

	Net Interest Income					
	(1)	(2)	(3)	(4)	(5)	(6)
$QE1_t \times Treat_i^Q$	0.051*** (0.014)	0.061*** (0.011)				
$QE2_t \times Treat_i^Q$	0.045** (0.019)	0.050*** (0.013)				
$QE3_t \times Treat_i^Q$	0.076** (0.032)	0.038** (0.017)				
$QE1_t \times Treat_i^D$			0.082*** (0.015)	0.074*** (0.009)		
$QE2_t \times Treat_i^D$			0.058*** (0.021)	0.048*** (0.013)		
$QE3_t \times Treat_i^D$			0.078*** (0.014)	0.019* (0.010)		
$QE1_t \times \left(\frac{MBS}{TotalAssets}\right)_i$					0.340*** (0.068)	0.354*** (0.056)
$QE2_t \times \left(\frac{MBS}{TotalAssets}\right)_i$					0.238** (0.100)	0.271*** (0.072)
$QE3_t \times \left(\frac{MBS}{TotalAssets}\right)_i$					0.467*** (0.174)	0.253* (0.136)

(continued)

Table 5. (Continued)

	Net Interest Income					
	(1)	(2)	(3)	(4)	(5)	(6)
Observations	12,786	11,040	5,148	4,445	24,997	21,526
R-squared	0.879	0.945	0.889	0.956	0.881	0.946
Bank-Level Controls	No	Yes	No	Yes	No	Yes
Bank Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Note: The dependent variable in columns 1–6 is logarithm of net interest income. $Treat_i^Q$ takes the value 1 for banks in the top quartile (top 25 percent) of the MBS-to-total-assets distribution, while 0 for banks in the bottom quartile (bottom 25 percent) of the distribution. $Treat_i^D$ takes the value 1 for banks in the top decile (top 10 percent) of the MBS-to-total-assets ratio, and 0 for banks in the bottom decile (bottom 10 percent) of the distribution. $\left(\frac{MBS}{TotalAssets}\right)_i$ is the ratio of MBS to total assets in 2007:Q4. $QE1_t$, $QE2_t$, and $QE3_t$ are dummy variables corresponding to the respective QE waves, where the dummy takes the value 1 during the QE episode and 0 otherwise. Control variables include bank size, liquidity, tier 1 capital ratio, and leverage ratio. All control variables are taken as lagged. Including interactions between controls and the treatment variables does not alter results significantly. The null hypothesis of cross-sectional independence in Pesaran’s (2021) test can be rejected for regression residuals at conventional significance levels. However, the average and average absolute cross-sectional correlations are relatively low, each standing approximately at 0.1 across regressions. According to Pesaran (2021), in finite samples, the assumption of cross-sectional independence is appropriate when cross-sectional correlations are low. Constant terms included but not reported. Robust standard errors in parentheses. ***, **, and * represent significance at 1 percent, 5 percent, and 10 percent, respectively.

are the interaction terms between the QE time dummies and banks' treatment status.

Our results overall suggest that treated banks generated larger profits during all episodes of QE. Specifically, in our sample of QE3, banks above the 75th percentile of the MBS-to-total-assets distribution increase their net interest income by about 3.8 percent (column 2) relative to banks below the 25th percentile of the MBS distribution. The results hold for different treatment definitions, i.e., based on the 10th and 90th percentiles of the MBS holdings (columns 3 and 4) and the continuous measure (columns 5 and 6). Our findings are consistent with the previous literature, including Flannery (1981), Bourke (1989), and Saunders and Schumacher (2000), that provides evidence of a positive association between monetary policy and bank profitability.

4.2.3 Systemic Risk

Our main empirical specification uses the Acharya et al. (2017) systemic expected shortfall index (*SES*) for bank i in year-quarter t relative to year-quarter $t-1$ as the dependent variable. The results are presented in Table 6. Again, we employ three treatment variables that classify banks based on quartiles, deciles, and the continuous measure of MBS to total assets.

Our main variable of interest, the interaction term $QE3_t \times Treat_i$, is negative and statistically significant across all definitions of the treatment variable, suggesting that banks with large MBS-to-total-assets holdings reduced their contribution to systemic risk during the third round of quantitative easing. The findings imply that QE had a positive effect on banking-sector stability. Our results are consistent with those of Berger, Roman, and Sedunov (2017), who report a positive impact of QE on the financial sector after the implementation of TARP.

Together, the results in Tables 4, 5, and 6 suggest that banks carrying higher MBS-to-total-assets ratios were characterized by higher levels of risk-taking and corresponding profitability in QE3, with evidence of higher profits in QE1 and QE2 as well. Finally, since the QE programs were largely unanticipated, especially the third round, banks that held more MBS had a prompt recovery and reduced their contribution to systemic risk.

Table 6. The Impact of QE on Systemic Risk

	Systemic Expected Shortfall					
	(1)	(2)	(3)	(4)	(5)	(6)
$QE1_t \times Treat_i^Q$	-0.014 (0.029)	-0.015 (0.029)				
$QE2_t \times Treat_i^Q$	-0.013 (0.046)	-0.011 (0.046)				
$QE3_t \times Treat_i^Q$	-0.092** (0.039)	-0.125*** (0.041)				
$QE1_t \times Treat_i^D$			-0.027 (0.055)	-0.027 (0.055)		
$QE2_t \times Treat_i^D$			-0.028 (0.084)	-0.015 (0.085)		
$QE3_t \times Treat_i^D$			-0.141* (0.080)	-0.134* (0.084)		
$QE1_t \times \left(\frac{MBS}{TotalAssets} \right)_i$					-0.100 (0.139)	-0.105 (0.136)
$QE2_t \times \left(\frac{MBS}{TotalAssets} \right)_i$					-0.056 (0.217)	-0.037 (0.213)
$QE3_t \times \left(\frac{MBS}{TotalAssets} \right)_i$					-0.363* (0.191)	-0.494** (0.192)

(continued)

Table 6. (Continued)

	Systemic Expected Shortfall					
	(1)	(2)	(3)	(4)	(5)	(6)
Observations	1,958	1,919	786	781	3,813	3,736
R-squared	0.193	0.175	0.125	0.128	0.209	0.217
Bank-Level Controls	No	Yes	No	Yes	No	Yes
Bank Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Note: The dependent variable in columns 1–6 is the change in SES_t . $Treat_i^Q$ takes the value 1 for banks in the top quartile (top 25 percent) of the MBS-to-total-assets distribution, while 0 for banks in the bottom quartile (bottom 25 percent) of the distribution. $Treat_i^D$ takes the value 1 for banks in the top decile (top 10 percent) of the MBS-to-total-assets ratio, and 0 for banks in the bottom decile (bottom 10 percent) of the distribution. $\left(\frac{MBS}{TotalAssets}\right)_i$ is the ratio of MBS to total assets in 2007:Q4. $QE1_t$, $QE2_t$, and $QE3_t$ are dummy variables corresponding to the respective QE waves, where the dummy takes the value 1 during the QE episode and 0 otherwise. Control variables include bank size, liquidity, tier 1 capital ratio, deposit ratio, and leverage ratio. All control variables are taken as lagged. Including interactions between controls and the treatment variables does not alter results significantly. The null hypothesis of cross-sectional independence in Pesaran’s (2021) test can be rejected for regression residuals at conventional significance levels. However, the average and average absolute cross-sectional correlations are relatively low, each standing approximately at 0.06 across regressions. According to Pesaran (2021), in finite samples, the assumption of cross-sectional independence is appropriate when cross-sectional correlations are low. Constant terms included but not reported. Robust standard errors in parentheses. ***, **, and * represent significance at 1 percent, 5 percent, and 10 percent, respectively.

4.3 *Systems Approach*

We also estimate our three specifications as a system of pooled equations via the seemingly unrelated regressions (SUR) model that employs the feasible generalized least squares (FGLS) estimator. A second system is estimated using the general method of moments (GMM) estimator. The system approach is adopted in order to account for potential cross-equation correlations in residuals and improve efficiency.²⁴ Possible endogeneity issues are further addressed via lags of variables as instruments in the GMM case. As is evident from Tables 7 and 8, the system estimates are consistent with the individual regression estimates in Tables 4–6.

4.4 *Too Big to Fail and Systemic Risk*

Following the GFC, the authorities attempted to prevent the failure of a few systemically important financial institutions (SIFIs) in order to avoid the risks posed to the overall financial system and broader economy. Failure of SIFIs, although uncommon, can have negative consequences for the entire financial system. These entities are often referred to as “too big to fail” (TBTF) due to their interconnected, sufficiently large, complex operations.

In the face of systemic failure, large financial institutions such as TBTFs accept government support, more commonly in the form of bailouts, that assists them in the recovery process (Cetorelli and Traina 2018). Accordingly, we focus on the impact of QE on systemic risk contributions by TBTF banks. This analysis draws motivation from the growing need to identify and monitor TBTFs, so that timely steps can be taken in order to avert events of economic distress associated with the failure of these systemically important banks.

We follow Zhou (2010) and identify the TBTF banks in our data.²⁵ To investigate the impact of QE on systemic risk

²⁴We analyze frameworks with control factors that differ across equations.

²⁵We are able to label 16 banks as TBTF. TBTF banks must satisfy the following selection criteria: (i) the financial institutions should be classified as “banks,” (ii) they should be traded on the NYSE, and (iii) they should be active from the beginning of 1987 until 2009.

Table 7. The Impact of QE on Bank Risk-Taking, Profitability, and Systemic Risk: Seemingly Unrelated Regressions

	Z-score	NII	SES
	(1)	(2)	(3)
$QE1_t \times Treat_i^Q$	-0.044 (0.122)	0.084 (0.067)	-0.324 (0.616)
$QE2_t \times Treat_i^Q$	0.003 (0.198)	0.069 (0.109)	-0.056 (0.051)
$QE3_t \times Treat_i^Q$	-0.529*** (0.140)	0.159** (0.077)	-1.56** (0.706)
Observations	1,928	1,928	1,928
R-squared	0.816	0.830	0.107
QE_t	Yes	Yes	Yes
Treatment Variable	Yes	Yes	Yes

Note: Pooled regressions employed. The dependent variable in column 1 is the logarithm of Z-score, in column 2 is the logarithm of net interest income (NII), and in column 3 is the systemic expected shortfall (SES). $Treat_i^Q$ takes the value 1 for banks in the top quartile (top 25 percent) of the MBS-to-total-assets distribution, while 0 for banks in the bottom quartile (bottom 25 percent) of the distribution. $QE1_t$, $QE2_t$, and $QE3_t$ are dummy variables corresponding to the respective QE waves, where the dummy takes the value 1 during the QE episode and 0 otherwise. Control variables include bank size, liquidity, tier 1 capital ratio, deposit ratio, and leverage ratio. All control variables are taken as lagged. Including interactions between controls and the treatment variables does not alter results significantly. The null hypothesis of cross-sectional independence in Pesaran’s (2021) test can be rejected for regression residuals at conventional significance levels. However, the average and average absolute cross-sectional correlations are relatively low, standing approximately at 0.06 and 0.06 for Z-scores; 0.04 and 0.05 for NII; and 0.03 and 0.05 for SES. According to Pesaran (2021), in finite samples, the assumption of cross-sectional independence is appropriate when cross-sectional correlations are low. Constant terms included but not reported. Robust standard errors in parentheses. ***, **, and * represent significance at 1 percent, 5 percent, and 10 percent, respectively.

contributions by TBTF banks, we estimate the following regression model:

$$\begin{aligned}
 Y_{i,t} = & \alpha_i + \beta_t + \rho_{i,t} + \gamma_1'QE_t + \gamma_2TBTF_i + \theta'TBTF_i \\
 & \times QE_t + \delta'X_{i,t-1} + \epsilon_{i,t},
 \end{aligned}
 \tag{6}$$

Table 8. The Impact of QE on Bank Risk-Taking, Profitability, and Systemic Risk: System GMM

	Z-score	NII	SES
	(1)	(2)	(3)
$QE1_t \times Treat_t^Q$	-0.097 (0.122)	0.171** (0.083)	-0.067* (0.039)
$QE2_t \times Treat_t^Q$	-0.037 (0.133)	0.256** (0.118)	-0.160** (0.065)
$QE3_t \times Treat_t^Q$	-0.273* (0.142)	0.567*** (0.191)	-0.816*** (0.276)
Observations	1,691	1,691	1,691
QE_t	Yes	Yes	Yes
Treatment Variable	Yes	Yes	Yes

Note: Pooled regressions employed. The dependent variable in column 1 is the logarithm of Z-score, in column 2 is the logarithm of net interest income (NII), and in column 3 is the systemic expected shortfall (SES). $Treat_t^Q$ takes the value 1 for banks in the top quartile (top 25 percent) of the MBS-to-total-assets distribution, while 0 for banks in the bottom quartile (bottom 25 percent) of the distribution. $QE1_t$, $QE2_t$, and $QE3_t$ are dummy variables corresponding to the respective QE waves, where the dummy takes the value 1 during the QE episode and 0 otherwise. Control variables include bank size, liquidity, tier 1 capital ratio, deposit ratio, and leverage ratio. All control variables are taken as lagged. Including interactions between controls and the treatment variables does not alter results significantly. The null hypothesis of cross-sectional independence in Pesaran's (2021) test can be rejected for regression residuals at conventional significance levels. However, the average and average absolute cross-sectional correlations are relatively low, standing approximately at 0.00 and 0.00 for Z-scores; 0.05 and 0.06 for NII; and 0.00 and 0.00 for SES. According to Pesaran (2021), in finite samples, the assumption of cross-sectional independence is appropriate when cross-sectional correlations are low. Constant terms included but not reported. Robust standard errors in parentheses. ***, **, and * represent significance at 1 percent, 5 percent, and 10 percent, respectively.

where $TBTF_i$ is a dummy variable that takes the value of 1 if a bank is identified as “too big to fail” and 0 otherwise. All other variables carry the same interpretation as in Equation (5) and the results are presented in Table 9.

TBTF banks are prone to moral hazard. If TBTF firms believe that the government will protect them during periods of distress, they have less incentive to monitor their financial operations (Dam


Table 9. The Impact of QE on Systemic Risk for Too-Big-to-Fail Banks

	Systemic Expected Shortfall	
	(1)	(2)
$QE1_t \times TBTF_i$	-0.008 (0.015)	-0.006 (0.014)
$QE2_t \times TBTF_i$	-0.080* (0.043)	-0.081* (0.045)
$QE3_t \times TBTF_i$	-0.138** (0.061)	-0.147** (0.064)
Observations	4,548	4,355
R-squared	0.235	0.220
Number of Banks	277	246
QE_t	Yes	Yes
Bank-Level Controls	No	Yes
Year-Quarter Fixed Effects	Yes	Yes
Bank Fixed Effects	Yes	Yes
State Fixed Effects	Yes	Yes

Note: The dependent variable is the growth rate of SES. $TBTF_i$ is a dummy that takes the value 1 for banks that are identified as systemically important financial institutions (SIFIs), or too big to fail (TBTF), by the Financial Stability Board and 0 otherwise. $QE1_t$, $QE2_t$, and $QE3_t$ are dummy variables corresponding to the respective QE waves, where the dummy takes the value 1 during the QE episode and 0 otherwise. Control variables include bank size, liquidity, tier 1 capital ratio, deposit ratio, and leverage ratio. All control variables are taken as lagged. Including interactions between controls and the treatment variables does not alter results significantly. The null hypothesis of cross-sectional independence in Pesaran's (2021) test can be rejected for regression residuals at conventional significance levels. However, the average and average absolute cross-sectional correlations are relatively low, standing approximately at 0.04 and 0.2. According to Pesaran (2021), in finite samples, the assumption of cross-sectional independence is appropriate when cross-sectional correlations are low. Constant terms included but not reported. Robust standard errors in parentheses. ***, **, and * represent significance at 1 percent, 5 percent, and 10 percent, respectively.

and Koetter 2012). With regards to the effectiveness of QE in reducing the contribution of TBTF banks to systemic risk, our results show that the policy had stabilizing effects. Moreover, the coefficient on the interaction term $QE3_t \times TBTF_i$ is -0.147 (column 2),

Figure 3. Robustness Test: Treatment Variable Based on Median



Note: The figure shows coefficient estimates where dependent variables are the logarithm of Z-score, the logarithm of net interest income, and the change in systemic expected shortfall. The regression uses an alternative treatment definition, with the dummy now equal to 1 if banks have an above-the-median ratio of MBS to total assets in 2007:Q4, and 0 if the ratio is below the median. Control variables include bank size, liquidity, tier 1 capital ratio, deposit ratio, and leverage ratio. All control variables are taken as lagged. The estimation includes bank fixed effects, year-quarter fixed effects, state effects, and lagged values of bank-level controls.

which is similar to the size of the effect obtained for the full sample of banks as discussed in the previous sections.

4.5 Other Robustness Checks

We perform a series of further robustness checks on our main results. First, we introduce a new treatment variable based on the median ratio of MBS holdings to total assets. This dummy variable takes the value of 1 if a bank is in the top half of the distribution of MBS to total assets in 2007:Q4 and 0 if it lies in the bottom half of the distribution. Figure 3 shows a plot of estimated coefficients θ

from Equation (5). Similar to the findings observed so far, it confirms that treated banks had significantly increased their capacity of risk-taking and profitability, and decreased their contribution to systemic risk during QE3.

Second, we use a bank's MBS exposure relative to securities, changing the specification to

$$Y_{i,t} = \alpha_i + \beta_t + \rho_{i,t} + \gamma'_1 QE_t + \gamma_2 \left(\frac{MBS}{Securities} \right)_i + \theta' \left(\frac{MBS}{Securities} \right)_i \times QE_t + \delta' X_{i,t-1} + \epsilon_{i,t}. \quad (7)$$

This measure has been employed in earlier studies, including that of Kurtzman, Luck, and Zimmermann (2022), and captures the treatment status of banks in relation to only the total securities, thus reducing any noise in the treatment specification when compared to the use of total assets in the denominator of the ratio. The results are presented in Table 10 and are qualitatively similar to the ones obtained in our main specifications. We still find stronger support for a differential increase in risk-taking and profitability, and a simultaneous decrease in systemic risk during QE3.

Third, we employ alternative measures of our dependent variables. We capture bank riskiness with the ratio of risk assets to total assets, while we use return on assets as a proxy for bank profitability. Systemic risk is captured by *SRISK*, which measures the amount of capital a bank would require for the maintenance of a given capital-asset ratio during periods of distress. *SRISK* is a function of bank size, which is captured by equity, leverage, and long-run *MES*. Following Brownlees and Engle (2017), *SRISK* is defined as

$$\begin{aligned} SRISK_{i,t} &= E_t[CapitalShortfall_{i,t+1}|Crisis] \\ &= E_t[k(Debt_{i,t+1} + Equity_{i,t+1}) - Equity_{i,t+1}|Crisis] \\ &= kDebt_{i,t} - (1-k)(1-LRMES_{i,t})Equity_{i,t}, \end{aligned} \quad (8)$$

where $Debt_{i,t}$ is the book value of debt, $Equity_{i,t}$ is the market value of equity, and k is the prudential capital fraction, i.e., level of book equity relative to assets, and is taken as 8 percent in line

Table 10. The Impact of QE on Bank Risk-Taking, Profitability, and Systemic Risk: Varying Treatment Definition

	Z-score		Net Interest Income		SES	
	(1)	(2)	(3)	(4)	(5)	(6)
$QE1_t \times Treat/Sec_i^Q$	0.005 (0.018)	0.014 (0.017)	0.038*** (0.008)	0.043*** (0.006)	-0.015 (0.022)	-0.016 (0.021)
$QE2_t \times Treat/Sec_i^Q$	-0.038* (0.022)	-0.031 (0.019)	0.041*** (0.011)	0.050*** (0.008)	0.001 (0.035)	-0.000 (0.034)
$QE3_t \times Treat/Sec_i^Q$	-0.095*** (0.020)	-0.071*** (0.021)	0.068*** (0.011)	0.043*** (0.008)	-0.065** (0.030)	-0.077** (0.031)
Observations	10,757	9,607	12,109	10,459	3,687	3,613
R-squared	0.646	0.665	0.984	0.993	0.229	0.236
Bank-Level Controls	No	Yes	No	Yes	No	Yes
Bank Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Note: The dependent variable in columns 1 and 2 is the logarithm of Z-score, in columns 3 and 4 is the logarithm of net interest income, and in columns 5 and 6 is the systemic expected shortfall (SES). $Treat/Sec_i^Q$ takes the value 1 for banks in the top quartile (top 25 percent) of the MBS-to-securities distribution, while 0 for banks in the bottom quartile (bottom 25 percent) of the distribution. $QE1_t$, $QE2_t$, and $QE3_t$ are dummy variables corresponding to the respective QE waves, where the dummy takes the value 1 during the QE episode and 0 otherwise. Control variables include bank size, liquidity, tier 1 capital ratio, deposit ratio, and leverage ratio. All control variables are taken as lagged. Including interactions between controls and the treatment variables does not alter results significantly. The null hypothesis of cross-sectional independence in Pesaran's (2021) test can be rejected for regression residuals at conventional significance levels. However, the average and average absolute cross-sectional correlations are relatively low, standing approximately at 0.07 and 0.07 for Z-scores; 0.07 and 0.07 for NI; and 0.03 and 0.05 for SES. According to Pesaran (2021), in finite samples, the assumption of cross-sectional independence is appropriate when cross-sectional correlations are low. Constant terms included but not reported. Robust standard errors in parentheses. ***, **, and * represent significance at 1 percent, 5 percent, and 10 percent, respectively.

with Brownlees and Engle (2017).²⁶ Finally, $LRMES_{i,t}$ is the long-run MES for bank i at time t , which is the average of returns on bank equity in the crisis period. Acharya, Engle, and Richardson (2012) approximate $LRMES_{i,t}$ as $1 - \exp\{-18 \times MES_{i,t}\}$. Table 11 shows that the coefficient on the interaction term between QE3 and the treatment variable remains significant and of the expected sign across dependent variables.

5. Conclusions

Bank risk-taking has gained much attention in recent times, as it is widely viewed as one of the outlets of unconventional monetary policy. The risk-taking channel also explains the reasons for the severity of the Global Financial Crisis. While individual risk-taking has been discussed in the literature, the impact on systemic risk is still less explored. Our paper explores the distributional effects of QE by examining whether banks that benefited more from the Fed's three rounds of QE also lowered their contribution to systemic risk.

We investigate whether QE-exposed banks increased their risk-taking capacity in order to achieve higher incomes and maximize profits. Since high incomes reflect the capability of a bank to generate higher profits from its asset management functions, we expect higher-income banks to be more stable. Furthermore, profitable banks have the necessary cushion to absorb losses, which reduces the probability of financial distress or failure.

Our results suggest that banks were encouraged to increase their risk-taking capacity and profitability due to the excess liquidity created by loose monetary policy. At the same time, banks reduced their contribution to systemic risk, suggesting that the implementation of QE had an overall positive effect on banking-sector stability. This is an important finding, as it shows that QE not only encouraged banks to take on higher risk, which led to increased availability of credit and liquidity, but also assisted in pulling the banking system out

²⁶The capital shortfall can be interpreted as the negative of working capital, i.e., the institution has a capital surplus when the capital shortfall is negative and suffers distress when capital shortfall is positive. Equation (8) assumes that debt cannot be renegotiated in the case of a crisis or systemic event.

Table 11. Varying Measures of Dependent Variable

	Risk/TA	ROA	SRISK
	(1)	(2)	(3)
$QE1_t \times Treat_i^Q$	-0.005*** (0.002)	0.001*** (0.000)	0.174 (0.346)
$QE2_t \times Treat_i^Q$	0.010*** (0.002)	0.001*** (0.000)	-0.782 (0.550)
$QE3_t \times Treat_i^Q$	0.013*** (0.002)	0.001*** (0.000)	-0.684* (0.405)
Observations	11,040	10,585	1,940
R-squared	0.758	0.654	0.903
Bank-Level Controls	Yes	Yes	Yes
Bank Fixed Effects	Yes	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes

Note: The dependent variable in column 1 is risk assets scaled by total assets, in column 2 is the return on assets, and in column 3 is *SRISK* scaled by bank equity. $Treat_i^Q$ takes the value 1 for banks in the top quartile (top 25 percent) of the MBS-to-total-assets distribution, while 0 for banks in the bottom quartile (bottom 25 percent) of the distribution. $QE1_t$, $QE2_t$, and $QE3_t$ are dummy variables corresponding to the respective QE waves, where the dummy takes the value 1 during the QE episode and 0 otherwise. Control variables include bank size, liquidity, tier 1 capital ratio, and leverage ratio. All control variables are taken as lagged. Including interactions between controls and the treatment variables does not alter results significantly. The null hypothesis of cross-sectional independence in Pesaran's (2021) test can be rejected for regression residuals at conventional significance levels. However, the average and average absolute cross-sectional correlations are relatively low, standing approximately at 0.07 and 0.07 for risk assets to total assets; 0.09 and 0.09 for ROA; and 0.04 and 0.05 for *SRISK*. According to Pesaran (2021), in finite samples, the assumption of cross-sectional independence is appropriate when cross-sectional correlations are low. Constant terms included but not reported. Robust standard errors in parentheses. ***, **, and * represent significance at 1 percent, 5 percent, and 10 percent, respectively.

of recession. Overall, our paper illustrates implications for individual banks and the financial system when assessing the distributional impact of QE.

We contribute to the literature by providing one of the first studies to investigate the impact of quantitative easing on the contribution of banks to systemic risk. Moreover, we investigate a potential mechanism that explains how QE encouraged banks to take up more risk in search of higher profits, thus making the banking sector more stable.

Appendix. Variables Employed: Construction and Corresponding Definitions

Table A.1. Variable Details

Variable Name	Definition	Data Sources
Securities Holdings	Held-to-maturity securities (BHCK1754) + available-for-sale securities (BHCK1773)	FR-Y9C
Treasury Securities	Trading Assets: Treasury Securities (BHCK3531)	FR-Y9C
Bank Size	Log of total assets (BHCK2170)	FR-Y9C
Equity Ratio	Total equity capital (BHCK3210) divided by total assets (BHCK2170)	FR-Y9C
Deposit Ratio	Non-interest-bearing deposits in domestic offices (BHDM6631) + interest-bearing deposits in domestic offices (BHDM6636) + non-interest-bearing deposits in foreign offices (BHFN6631) divided by total assets (BHCK2170)	FR-Y9C
Liquidity	Cash and balances due from depository institutions: non-interest-bearing balances and currency and coin (BHCK0081) + interest-bearing balances in U.S. offices (BHCK0395) + interest-bearing balances in foreign offices, Edge and Agreement subsidiaries, and IBFs (BHCK0397) divided by total assets (BHCK2170)	FR-Y9C
Total Lending	Total loans (BHCK2122) divided by total assets (BHCK2170)	FR-Y9C
Net Interest Income	Logarithm of net interest income (BHCK4074)	FR-Y9C
Return on Assets	Net income (BHCK4340) divided by average total assets (BHCT3368)	FR-Y9C
Z-score	Sum of ROA and equity ratio divided by standard deviation of asset returns	FR-Y9C
Risk Assets	Total assets (BHCK2170) – cash and balances due from depository institutions (BHCK0081 + BHCK0395 + BHCK0397) – federal funds sold (BHDMB987) – securities purchased under agreement to resell (BHCKB989) divided by total assets (BHCK2170)	FR-Y9C
SES	Propensity of a bank to be undercapitalized when the entire banking system is undercapitalized	CRSP and FR-Y9C
SRISK	Amount of capital that is required by a financial institution during periods of crisis	CRSP and FR-Y9C

Note: The table presents data sources and method of construction of variables used in analysis. FR-Y9C refers to balance sheet information of all BHCs from the Federal Reserve Bank of Chicago. CRSP refers to market data in the Center for Research in Security Prices database.

References

- Abadi, J., M. Brunnermeier, and Y. Koby. 2023. "The Reversal Interest Rate." *American Economic Review* 113 (8): 2084–2120.
- Abbate, A., and D. Thaler. 2019. "Monetary Policy and the Asset Risk-Taking Channel." *Journal of Money, Credit and Banking* 51 (8): 2115–44.
- Acharya, V., R. Engle, and M. Richardson, eds. 2012. "Capital Shortfall: A New Approach to Ranking and Regulating Systemic Risks." *American Economic Review* 102 (3): 59–64.
- Acharya, V. V., L. H. Pedersen, T. Philippon, and M. Richardson. 2017. "Measuring Systemic Risk." *Review of Financial Studies* 30 (1): 2–47.
- Acharya, V. V., and M. P. Richardson. 2009. *Restoring Financial Stability: How to Repair a Failed System*. John Wiley and Sons.
- Acharya, V. V., and T. Yorulmazer. 2007a. "Cash-in-the-Market Pricing and Optimal Resolution of Bank Failures." *Review of Financial Studies* 21 (6): 2705–42.
- . 2007b. "Too Many to Fail—An Analysis of Time-Inconsistency in Bank Closure Policies." *Journal of Financial Intermediation* 16 (1): 1–31.
- Adrian, T., and M. K. Brunnermeier. 2011. "CoVaR." NBER Working Paper No. 17454.
- Alessandri, P., and B. D. Nelson. 2015. "Simple Banking: Profitability and the Yield Curve." *Journal of Money, Credit and Banking* 47 (1): 143–75.
- Allen, L., T. G. Bali, and Y. Tang. 2012. "Does Systemic Risk in the Financial Sector Predict Future Economic Downturns?" *Review of Financial Studies* 25 (10): 3000–3036.
- Altunbas, Y., L. Gambacorta, and D. Marques-Ibanez. 2010. "Does Monetary Policy Affect Bank Risk-Taking?" BIS Working Paper No. 298.
- Beck, T., O. De Jonghe, and G. Schepens. 2013. "Bank Competition and Stability: Cross-Country Heterogeneity." *Journal of Financial Intermediation* 22 (2): 218–44.
- Berger, A. N., R. A. Roman, and J. Sedunov. 2017. "Do Bank Bailouts Reduce or Increase Systemic Risk? The Effects of TARP on Financial System Stability."

- Bernanke, B. S. 2007. "The Financial Accelerator and the Credit Channel." Speech at The Credit Channel of Monetary Policy in the Twenty-first Century Conference, Federal Reserve Bank of Atlanta, Atlanta, Georgia, June 15.
- Bernanke, B. S., T. F. Geithner, J. N. Liang, and H. M. Paulson. 2020. *First Responders: Inside the US Strategy for Fighting the 2007-2009 Global Financial Crisis*. Yale University Press.
- Bernanke, B. S., and M. Gertler. 1995. "Inside the Black Box: The Credit Channel of Monetary Policy Transmission." *Journal of Economic Perspectives* 9 (4): 27–48.
- Bikker, J. A., and T. M. Vervliet. 2018. "Bank Profitability and Risk-Taking under Low Interest Rates." *International Journal of Finance and Economics* 23 (1): 3–18.
- Borio, C., L. Gambacorta, and B. Hofmann. 2017. "The Influence of Monetary Policy on Bank Profitability." *International Finance* 20 (1): 48–63.
- Borio, C., and H. Zhu. 2012. "Capital Regulation, Risk-Taking and Monetary Policy: A Missing Link in the Transmission Mechanism?" *Journal of Financial Stability* 8 (4): 236–51.
- Bourke, P. 1989. "Concentration and Other Determinants of Bank Profitability in Europe, North America and Australia." *Journal of Banking and Finance* 13 (1): 65–79.
- Brownlees, C., and R. F. Engle. 2017. "SRISK: A Conditional Capital Shortfall Measure of Systemic Risk." *Review of Financial Studies* 30 (1): 48–79.
- Brunnermeier, M. K., and Y. Koby. 2016. "The Reversal Interest Rate: An Effective Lower Bound on Monetary Policy." Unpublished Paper, Princeton University.
- Cetorelli, N., and J. Traina. 2018. "Resolving 'Too Big to Fail'." Staff Report No. 859, Federal Reserve Bank of New York.
- Chakraborty, I., I. Goldstein, and A. MacKinlay. 2020. "Monetary Stimulus and Bank Lending." *Journal of Financial Economics* 136 (1): 189–218.
- Choi, D. B. 2014. "Heterogeneity and Stability: Bolster the Strong, Not the Weak." *Review of Financial Studies* 27 (6): 1830–67.
- Cordella, T., and E. L. Yeyati. 2003. "Bank Bailouts: Moral Hazard vs. Value Effect." *Journal of Financial Intermediation* 12 (4): 300–330.

- Dam, L., and M. Koetter. 2012. "Bank Bailouts and Moral Hazard: Evidence from Germany." *Review of Financial Studies* 25 (8): 2343–80.
- Delis, M. D., and G. P. Kouretas. 2011. "Interest Rates and Bank Risk-Taking." *Journal of Banking and Finance* 35 (4): 840–55.
- Dell’Ariccia, G., L. Laeven, and R. Marquez. 2014. "Real Interest Rates, Leverage, and Bank Risk-Taking." *Journal of Economic Theory* 149 (January): 65–99.
- Dell’Ariccia, G., and L. Ratnovski. 2013. "Bailouts and Systemic Insurance." IMF Working Paper No. 13/233.
- Diamond, D. W., and R. G. Rajan. 2005. "Liquidity Shortages and Banking Crises." *Journal of Finance* 60 (2): 615–47.
- . 2009a. "Illiquidity and Interest Rate Policy." Technical Report, National Bureau of Economic Research.
- . 2009b. "The Credit Crisis: Conjectures about Causes and Remedies." *American Economic Review* 99 (2): 606–10.
- Farhi, E., and J. Tirole. 2012. "Collective Moral Hazard, Maturity Mismatch, and Systemic Bailouts." *American Economic Review* 102 (1): 60–93.
- Financial Stability Board. 2009. "Report to G20 Finance Ministers and Governors Guidance to Assess the Systemic Importance of Financial Institutions, Markets and Instruments: Initial Considerations." *Wayamba Journal of Management* 10 (1).
- Flannery, M. J. 1981. "Market Interest Rates and Commercial Bank Profitability: An Empirical Investigation." *Journal of Finance* 36 (5): 1085–1101.
- Freixas, X., A. Martin, and D. Skeie. 2011. "Bank Liquidity, Interbank Markets, and Monetary Policy." *Review of Financial Studies* 24 (8): 2656–92.
- Gagnon, J., and M. Holscher. 2008. "Purchases of Agency MBS and Debt." Federal Open Market Committee (FOMC), Federal Reserve System.
- Galstyan, V., and A. Velic. 2017. "Debt Thresholds and Real Exchange Rates: An Emerging Markets Perspective." *Journal of International Money and Finance* 70 (February): 452–70.
- Gambacorta, L. 2009. "Monetary Policy and the Risk-Taking Channel." *BIS Quarterly Review* (December): 43–54.
- Huang, X., H. Zhou, and H. Zhu. 2012. "Systemic Risk Contributions." *Journal of Financial Services Research* 42 (1–2): 55–83.

- Ioannidou, V., S. Ongena, and J.-L. Peydró. 2014. "Monetary Policy, Risk-Taking, and Pricing: Evidence from a Quasi-natural Experiment." *Review of Finance* 19 (1): 95–144.
- Jiménez, G., S. Ongena, J.-L. Peydró, and J. Saurina. 2014. "Hazardous Times for Monetary Policy: What Do Twenty-Three Million Bank Loans Say about the Effects of Monetary Policy on Credit Risk-Taking?" *Econometrica* 82 (2): 463–505.
- Kandrac, J., and B. Schlusche. 2021. "Quantitative Easing and Bank Risk Taking: Evidence from Lending." *Journal of Money, Credit and Banking* 53 (4): 635–76.
- Kapoor, S., and O. Peia. 2021. "The Impact of Quantitative Easing on Liquidity Creation." *Journal of Banking and Finance* 122 (January): Article 105998.
- Keeley, M. C. 1990. "Deposit Insurance, Risk, and Market Power in Banking." *American Economic Review* 80 (5): 1183–1200.
- Krishnamurthy, A., and A. Vissing-Jorgensen. 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." NBER Working Paper No. 17555.
- . 2014. "The Ins and Outs of LSAPs." In *Global Dimensions of Unconventional Monetary Policy*. Proceedings of the 2013 Economic Policy Symposium. Kansas City, MO: Federal Reserve Bank of Kansas City.
- Kurtzman, R., S. Luck, and T. Zimmermann. 2022. "Did GE Lead Banks to Relax Their Lending Standards? Evidence from the Federal Reserve's LSAPs." *Journal of Banking and Finance* 138 (May): Article 105403.
- Laeven, L., and R. Levine. 2009. "Bank Governance, Regulation and Risk Taking." *Journal of Financial Economics* 93 (2): 259–75.
- Lee, C.-C., and M.-F. Hsieh. 2014. "Bank Reforms, Foreign Ownership, and Financial Stability." *Journal of International Money and Finance* 40 (February): 204–24.
- Luck, S., and T. Zimmermann. 2020. "Employment Effects of Unconventional Monetary Policy: Evidence from GE." *Journal of Financial Economics* 135 (3): 678–703.
- Martynova, N. 2015. "Bank Profitability and Risk-Taking." IMF Working Paper No. 15/249.
- Nijskens, R., and W. Wagner. 2011. "Credit Risk Transfer Activities and Systemic Risk: How Banks Became Less Risky Individually

- but Posed Greater Risks to the Financial System at the Same Time.” *Journal of Banking and Finance* 35 (6): 1391–98.
- Paligorova, T., and J. A. Santos. 2017. “Monetary Policy and Bank Risk-Taking: Evidence from the Corporate Loan Market.” *Journal of Financial Intermediation* 30 (April): 35–49.
- Pesaran, M. H. 2021. “General Diagnostic Tests for Cross-Sectional Dependence in Panels.” *Empirical Economics* 60 (1): 13–50.
- Rajan, R. G. 2006. “Has Finance Made the World Riskier?” *European Financial Management* 12 (4): 499–533.
- Rodnyansky, A., and O. M. Darmouni. 2017. “The Effects of Quantitative Easing on Bank Lending Behavior.” *Review of Financial Studies* 30 (11): 3858–87.
- Saunders, A., and L. Schumacher. 2000. “The Determinants of Bank Interest Rate Margins: An International Study.” *Journal of International Money and Finance* 19 (6): 813–32.
- Schaeck, K., and M. Čihák. 2010. “Competition, Efficiency, and Soundness in Banking: An Industrial Organization Perspective.” ECB Discussion Paper No. 2010-205.
- Velic, A. 2025. “Relative Finance Wages and Inequality: A Role for Intangibles?” *Journal of International Financial Markets, Institutions and Money* 103 (September): Article 102192.
- Williams, B. 2016. “The Impact of Non-interest Income on Bank Risk in Australia.” *Journal of Banking and Finance* 73 (December): 16–37.
- Zhou, C. 2010. “Are Banks Too Big to Fail? Measuring Systemic Importance of Financial Institutions.” *International Journal of Central Banking* 6 (4): 205–50.

Financial Stability Implications of CBDC*

Francesca Carapella, Jin-Wook Chang, Sebastian Infante,
Melissa Leistra, Arazi Lubis, and Alexandros P. Vardoulakis
Federal Reserve Board

A central bank digital currency (CBDC) is a form of digital money that is denominated in the national unit of account, constitutes a direct liability of the central bank, and can be distinguished from other central bank liabilities. We examine the financial stability risks and benefits of issuing a CBDC under different design options. Our analysis is based on lessons derived from historical case studies as well as on an analytical framework that allows us to characterize the mechanisms through which a CBDC can affect financial stability. We further discuss various policy tools that can be employed to mitigate financial stability risks.

JEL Codes: E40, E50, G01, G21, G23, G28.

1. Introduction

A central bank digital currency (CBDC) is a form of digital money that is denominated in the national unit of account, constitutes a direct liability of the central bank, and can be distinguished from other central bank liabilities.¹ This paper examines the financial stability implications of a CBDC under different design options and discusses policy tools that can be employed to mitigate financial

*The views expressed in this paper are those of the authors and do not necessarily represent those of the Federal Reserve Board or anyone in the Federal Reserve System. We thank Michael Kiley, Beth Klee, Andreas Lehnert, Daniel Sanches, Chiara Scotti, and Skander Van den Heuvel for useful comments and suggestions.

¹Most central banks already provide digital liabilities in the form of reserve balances, which are mainly accessible to banks. Common definitions of CBDC do not attempt to determine on which dimensions CBDC could or should vary from other digital central bank monies. There are many potential design and policy choice variations. See, for example, Committee on Payments and Market Infrastructure (2018) at <https://www.bis.org/cpmi/publ/d174.pdf>.

stability risks, abstracting from any other issues not related to financial stability. Because a CBDC could in concept be held or transferred by any party such as institutional or retail market participants, we focus our analysis on a general-purpose CBDC, though some of the channels we discuss may be more relevant for specific wholesale or retail use-cases.

We examine three broad ways in which a CBDC may affect financial stability.² First, a CBDC may increase the financial sector's vulnerability to destabilizing runs. By offering a safer asset, relative to private liabilities of the financial sector, a CBDC may present an attractive option for depositors and investors to fly to during times of stress. Moreover, a CBDC may be more attractive than other forms of government debt that are also safe but are more price sensitive during times of market stress or less useful as a direct means of payment.³ Therefore, because of CBDC's relative safety and convenience, some nonbank institutions, such as money funds, might suffer larger or more frequent runs in the presence of it. Policy tools, such as CBDC remuneration and quantity limits on CBDC holdings, have the potential to mitigate such run risks, complementing traditional central bank tools for emergency liquidity assistance. Tiered remuneration—whereby remuneration decreases as CBDC holdings increase and can be adjusted during times of stress—could be combined with a hard quantity limit on individual counterparty holdings. The holding limit could be chosen to prevent severe runs to CBDC, while payments in CBDC continue to be possible. The calibration of holding limits is an important but difficult task. Given the uncertainty about the demand for a CBDC, setting holding limits may involve a process of trial and error.

Second, a CBDC may weaken financial stability by reducing the ability of banks to extend credit during times of stress, which could result in a more widespread credit crunch. We study a stress scenario in which the presence of a CBDC causes a withdrawal of

²This paper focuses on financial stability risks from a CBDC during times of stress rather than on the potential effect of a CBDC on bank deposits and credit provision during normal times.

³Vissing-Jorgensen (2021) documents that the 10-year Treasury yield increased by 64 basis points from March 9 to March 18, 2020, making 10-year Treasury securities a particularly unattractive investment option during the market turmoil at the onset of the COVID-19 pandemic.

deposits from banks, forcing banks to use less desirable funding sources despite not necessarily precipitating a crisis. Using various estimates from the literature, we find that funding costs and, in turn, the rates charged to borrowers, may increase by 50 to 250 basis points in stress times. We roughly estimate that such higher borrowing rates may result in a decrease of about 1 to 5 percent in commercial and industrial (C&I) lending. Furthermore, the fear of deposit losses to CBDC in times of stress may make banks less willing to provide ex ante liquidity insurance to firms in the form of credit lines. Kashyap, Rajan, and Stein (2002) and Piazzesi and Schneider (2022) highlight the complementarity between deposit-taking and the extension of credit lines to firms, accruing from the fact that deposit withdrawals may not perfectly correlate with credit-line drawdowns. This imperfect correlation allows bank to economize on the holdings of liquid assets required to insure against the risk of outflows. To the extent that CBDC increases this correlation in times of stress, the provision of liquidity insurance to firms would decrease, further amplifying the adverse effects of a credit crunch.

Third, to the extent a CBDC is made more broadly available than other digital central bank liabilities, it could make the financial system more stable by providing a perfectly safe asset to a wider set of financial firms than those that currently make use of central bank liabilities, such as nonbanks or foreign institutions, which might use it to enhance their liquidity management. Beyond its effect on the traditional financial system, a CBDC could improve financial stability in the digital asset ecosystem. Stablecoins with direct access to a digital central bank liability would be better able to prevent or manage runs.

Our analysis proceeds as follows. Section 2 presents case studies that share similar financial stability concerns as the ones raised for CBDC. Section 3 discusses specific mechanisms that have been highlighted in economic research through which a CBDC may affect financial stability. Section 4 analyzes the various tools to mitigate financial stability risks introduced by a CBDC and their implementation challenges.⁴

⁴See Carapella and Flemming (2020), Infante et al. (2022, 2024), and Ahnert et al. (2024) for detailed literature reviews.

2. Lessons from Case Studies

Because a CBDC offers a safe asset to run to in times of crisis, a key financial stability concern about its introduction is that it may increase the vulnerability of institutions and markets that depend on short-term, runnable debt. Flight-to-safety concerns are not unique to CBDC and have been raised in several examples both in the U.S. and abroad. We review these cases to infer lessons that are relevant for thinking about any destabilizing effects of a flight to CBDC in times of crisis.

First, we review similar concerns introduced by potential requests for accounts at the Federal Reserve by state-chartered narrow banking institutions, planning to pass through part of the interest rates they would earn on their reserve holdings to their depositors. Second, we examine the similarities between a flight to CBDC and a flight to ON RRP (the Federal Reserve's overnight reverse repurchase facility), as an example of a central bank liability available to eligible nondepository institutions, and consider whether such facility has contributed to the run on prime money market funds (MMFs) during the COVID-19 crisis in March 2020.⁵ Third, we analyze the recent episodes of traditional bank runs, including the run on Greek banks during the European sovereign debt crisis and banking stress of March 2023.

2.1 *Pass-Through Investment Entities*

Pass-through investment entities (PTIEs), a type of “narrow bank,” were contemplated in a 2019 request for comment by the Federal Reserve Board (Board). The request for comment explained that such institutions potentially could seek to be granted master accounts at the Federal Reserve and to be eligible for receiving interest on reserves balances (IORB). Their business model would be to invest a significant fraction, or all, of their total deposits in reserve balances with the Federal Reserve in order to pass through the interest they earn on reserve balances to their depositors. The

⁵The ON RRP is a supplementary monetary policy tool that supports interest rate control by providing a floor in overnight money market rates to eligible investors.

Board's previous consideration of potential PTIEs provides the most direct comparison with the flight-to-safety concerns introduced by a CBDC.

If established and deemed eligible institutions, PTIEs would be businesses operating on narrow profit margins and, because their investments would be in perfectly safe assets, they potentially could be exempted from prudential regulation and paying insurance deposit premiums. As a result, there would be no restrictions, in principle, on PTIEs' size. To mitigate the potential risks arising from the attractiveness of a nearly unlimited safe asset in times of stress, among other reasons, the Board issued an Advanced Notice of Proposed Rulemaking (ANPR) to request public comment whether it should propose changes to Regulation D, such as on how to define PTIEs and on potentially paying them a lower remuneration rate than IORB (including possibly zero) or tiered rates for different levels of balances held by PTIEs.⁶

2.1.1 Lessons for CBDC

No PTIEs have been established so far, but we can still draw important lessons for CBDC. The appeal of PTIEs during times of stress would accrue both from the safety of their assets and the potential interest paid to their depositors. Paying a lower or zero rate of interest on reserve holdings of PTIEs—potentially above a certain holding threshold as contemplated by the ANPR—would lower the interest that PTIEs could offer depositors and potentially mitigate the incentives to run to PTIEs in times of stress. The proposal to amend Regulation D shares similar characteristics to tiered CBDC remuneration to mitigate a flight to CBDC, a policy tool we explore later.

2.2 ON RRP, Bank Deposits, and MMFs

Another arrangement that is comparable to CBDC for its property of being a potential safe haven during crisis times is the overnight

⁶The ANPR was published in *Federal Register* Vol. 84, No. 48 on Tuesday, March 12, 2019 (Federal Register :: Regulation D: Reserve Requirements of Depository Institutions). The Board had not determined whether any or all potential PTIEs might be eligible institutions.

reverse repurchase (ON RRP) facility. The Federal Reserve introduced the ON RRP in 2014 to support its control of short-term interest rates by allowing eligible nonbank participants in the short-term funding markets to earn interest on repurchase agreements (repos) with the Federal Reserve. The set of such eligible counterparties is broad and includes money market funds (MMFs), primary dealers, and government-sponsored enterprises (GSEs). Although the ON RRP is a monetary policy tool, its potential for destabilizing effects was highlighted early on.⁷ In theory, there are two channels through which access to the ON RRP could potentially affect financial stability.

The first channel operates through bank depositors who may run on their banks and move funds into government MMFs during a flight-to-safety episode.⁸ Large institutional entities have funds that they choose to keep in their banks or to invest in short-term instruments, balancing their needs for immediate liquidity, credit risk management, and potential returns. In times of stress, the incentive to move bank account balances that are above deposit insurance thresholds to secured investments increases. Government MMFs could receive many of these investments. Without access to ON RRP, government MMFs would invest these additional funds into government assets, such as Treasury securities, which would put downward pressure on their yields, mitigating the incentive to run. This may incentivize MMFs to take actions to limit additional inflows when the Treasury yield falls significantly. With access to an elastic ON RRP, instead, MMFs could absorb more liquidity without suppressing short-term rates. As a result, access to ON RRP could exacerbate potential runs from banks to MMFs.

The second channel operates through nonbank institutions, such as MMFs, that may substitute away from lending to financial or non-financial firms toward the ON RRP. MMFs are large cash lenders in repo markets, substantial buyers of commercial paper, and otherwise lend in short-term funding markets. A migration of such funds

⁷See, for example, Frost et al. (2015).

⁸Government MMFs are defined as MMFs that invest 99.5 percent or more of their total assets in very liquid investments, namely, cash, government securities, and/or repurchase agreements that are collateralized fully with government securities.

to ON RRP could amplify a crisis by reducing multiple sources of short-term funding for financial and nonfinancial firms. The financial stability risks of both channels are higher under full allotment, that is, if the facility is uncapped on aggregate, or if a cap is sufficiently large.⁹ To mitigate these risks, the ON RRP imposes individual caps, while the size of the System Open Market Account (SOMA) portfolio also imposes an implicit aggregate limit. For example, the initial limit on the amount each counterparty could invest was initially set to \$30 billion, was raised to \$80 billion in March 2021, and then raised to \$160 billion in September 2021.¹⁰


Stress events since the introduction of ON RRP are instructive for deriving potential lessons for a CBDC. We focus on the events at the outbreak of the COVID-19 pandemic when prime, in particular, MMFs experienced runs (see Li et al. 2021). As shown in Figure 1, the ON RRP take-up spiked in March–April 2020. However, flows to the ON RRP during that stress episode were limited overall, and take-up quickly normalized to levels close to zero. At the same time, deposits, including uninsured ones, increased at a fast rate.¹¹ Similarly, there is limited evidence that the ON RRP contributed to the banking stress of March of 2023. While some regional and small banks suffered a contraction in their deposit base, ON RRP take-up declined in the first week following the onset of the crisis. These episodes indicate that there is little evidence to date for the first channel described above, through which ON RRP can have destabilizing effects.¹²

⁹A full allotment facility is limited by the overall size of the Federal Reserve's balance sheet. In the highly unlikely event that the ON RRP demand exceeds the amount of available securities, awards will be made at the rate at which this size limit was achieved (the stop-out rate), with all propositions below this rate awarded in full and all propositions equal to this rate awarded on a pro rata basis. See https://www.newyorkfed.org/markets/rrp_faq.

¹⁰Initially the facility also had an exogenously set aggregate limit of \$300 billion, which was removed in December 2015. Before its removal, the aggregate cap was reached once in September 2014.

¹¹The increase in ON RRP take-up that started occurring in March 2021 is not related to flight-to-safety episodes and, thus, not relevant for financial stability.

¹²See Logan (2021) for a similar conclusion. We should note, however, that the March 2020 episode was primarily a liquidity event that found banks with strong capital and liquidity position. Movement from bank deposits to ON RRP amid stress driven by concerns about bank credit quality might look different.

Figure 1. ON RRP and Deposit Flows

Source: Federal Reserve Economic Data (FRED) and Call Reports.

Note: Data on deposits are at a quarterly frequency. The ON RRP data are daily and the last observation is November 6, 2023.


The second channel briefly emerged during the COVID episode but was not quantitatively strong nor persistent. As shown in the left panel of Figure 2, prime MMFs' holdings of certificate of deposits (CDs) and commercial paper (CP) fell while ON RRP take-up briefly increased along with Treasury holdings. Quickly, ON RRP take-up returned to previous levels while Treasury growth continued. The right panel of Figure 2 shows a similar pattern for government MMFs.

One possible explanation why the second channel, described above, was relatively weak was because there were liquid and safe alternatives, namely Treasury bills, yielding higher rates. While the ON RRP rate and the yield on four-week Treasury bills fell after the COVID-19 shock, Treasury-bill yields traded relatively higher than the ON RRP rate, and Treasuries could absorb the extra demand for safety (Figure 3). It is conceivable that the second channel could have been stronger if Treasury bills would have been more price inelastic during this period, pushing Treasury yields down, potentially to negative territory, and making ON RRP more attractive.


The ON RRP may have facilitated a modest portion of the outflows from prime MMFs and inflows toward government MMFs, but other factors may also have facilitated the reallocation from prime to government MMFs. In particular, relatively low switching costs for investors to move their funds out of prime MMF to other vehicles have been found to be important factors for run intensity.¹³

¹³See Cipriani and La Spada (2020).

Figure 2. Exposures of Prime MMFs (left chart) and Government MMFs (right chart)



Source: U.S. Securities and Exchange Commission (SEC); EDGAR, form N-MFP data.

Figure 3. ON RRP and Treasury Yields

Source: FRED.

Moreover, fund families with relatively higher specialization in government MMFs saw greater outflows from their prime MMFs into their government funds, both for institutional and retail investors.

2.2.1 Lessons for CBDC

There are at least three lessons helpful for thinking about financial stability risks from a CBDC. First, addressing the source of funding risks at the institutional level may be effective at mitigating financial stability concerns from a flight to safety. Prime MMFs experienced runs at the outbreak of the COVID-19 pandemic, but strong capital and liquidity positions of banks were likely the reasons why confidence in large banks was never seriously shaken and large uninsured depositors did not move their funds from large banks to government MMFs with access to the ON RRP. While the stress episode of March 2023 does suggest that some smaller and less creditworthy banks are subject to depositor runs, the role of the ON RRP in attracting them was also limited.¹⁴

Second, the rate of remuneration on a CBDC during crisis would be crucial. As discussed above, ON RRP take-up likely was limited

¹⁴The Federal Reserve issued an after-action report on bank and supervision shortcomings following the failure of Silicon Valley Bank. See <https://www.federalreserve.gov/publications/files/svb-review-20230428.pdf>.


during the pandemic because the ON RRP rate was set to zero, making it less attractive relative to other safe alternatives. Because a CBDC may offer benefits to holders that are not provided by other money market instruments or ON RRP balances, the remuneration rate on a CBDC may need to be considerably lower than in normal times to discourage a flight to safety during a crisis. Holding limits on CBDC may also be useful.

Third, as suggested by the experience of prime MMFs during March 2020, low switching costs may facilitate run episodes on banks and other institutions investing in money market instruments. One distinct difference between a CBDC and the ON RRP is that the former could serve as a settlement asset. The potential features of CBDC or its supporting infrastructure that could make it attractive, such as enabling fast payments, facilitating interoperability, or provision of wider access to digital central bank money, could dramatically reduce the switching costs from both deposits and money market instruments and lead to much faster runs.¹⁵

2.3 Recent Episodes of Traditional Bank Runs

Recently, there have been limited cases of traditional banking panics with system-wide implications. Two salient examples are the run on the Greek banking system during the European sovereign debt crisis and the banking stress in March 2023 in the United States. These episodes provide good examples of how a run may materialize at a slow pace—with high switching costs from bank deposits to other alternative stores of value, including physical currency—or in a matter of a days when switching costs are low and depositors are connected to one another. The first example suggests that imposing switching costs from deposits to a CBDC may slow the pace of a run, but it cannot eliminate runs altogether. The second example

¹⁵In the long run, low switching costs could reduce a bank's moral hazard, as the fear of a run could discipline the bank, urging it to adapt its business model appropriately (Calomiris and Kahn 1991). That said, the overall effect of low switching costs could lead to faster runs, as evidenced by the stress episode in March 2023 discussed in the next subsection. In addition, the reduction of switching costs can be interpreted as an improvement in interoperability—a potentially desired feature of CBDC—which can become a contagion channel in times of stress.

Figure 4. Cash and Deposit Holdings during Greek Crisis

Source: Bank of Greece, European Central Bank, <https://www.bankofgreece.gr/en/statistics/monetary-and-banking-statistics/monetary-aggregates>.

suggests that reduced switching costs and fast convertibility offered by a CBDC may have the potential to further exacerbate runs.

The Greek case materialized amid concerns over sovereign debt sustainability and the hypothetical exit of Greece from the euro zone. Deposits at Greek banks fell by just over 30 percent from 2010 to 2012 and by a further 25 percent from 2014 to 2016. In this instance, the deposit runoff was slow moving. Most Greek residents had to rely on domestic banks or physical banknotes to store their liquid wealth and to make and receive payments, while only larger firms and wealthier individuals already had or could easily establish bank accounts in other countries.¹⁶ Figure 4 plots the withdrawal of deposits along with the spikes in cash holdings in Greece during this period. Note that after the imposition of capital controls in June 2015, demand deposits started mechanically increasing and available cash started mechanically falling.

Interest paid on deposits may have been a reason why retail depositors kept their money at Greek banks. Another factor for the slow exodus of deposits from banks was likely the switching costs associated with withdrawing deposits to hold physical currency. In fact, storing cash in a safe deposit box at a bank requires a separate process and could have been expensive for small depositors, while it also entails “shoe-leather” costs every time a depositor wants to

¹⁶See, for example, Artavanis et al. (2019).

access the stored cash.¹⁷ Moreover, individuals who withdrew larger amounts of cash were a common target of house robberies and individual assaults.¹⁸ An additional switching cost from bank deposits to cash was the length of time needed to obtain physical currency from a bank. Depositors had to call their banks several days in advance so that the bank could bring in the required cash.

The more recent case of the banking stress of March 2023 shows a dramatically different environment, which resulted in deposit runs that were extraordinarily fast and large. Rose (2023) documents that the runs on Silicon Valley Bank and Signature Bank materialized in a matter of one or two days. In contrast to the Greek stress, confidence in other U.S. banks remained strong, and thus depositors willingly transferred deposits to other banks, which they were able to do quickly. Sectoral or other connections between bank depositors facilitated the particularly rapid and coordinated withdrawal. While the concentration of uninsured depositors was a particularly salient issue during this bank run episode, it underscores how the availability of a safe store of value with seamless and immediate convertibility can quickly destabilize affected banks. The introduction of a CBDC might similarly accelerate runs through any removal of payments frictions its transfer system entails and, perhaps more uniquely, through provision of a safe alternative that is not dependent on continued confidence in the wider banking sector.

2.3.1 Lessons for CBDC

While the parallels with the potential effects of a CBDC are imperfect, the Greek and March 2023 episodes illustrate how switching costs can affect run dynamics. In both cases, conditions prompted a desire to run from banks that were deemed illiquid or in danger of being insolvent. In the case of Greek banks, depositors that faced low switching costs to readily available alternatives left affected banks

¹⁷Historical data on the rental cost of safe deposit boxes are not readily available. The National Bank of Greece is currently offering them at the range of 100 to 500 euros. See https://www.nbg.gr/-/jssmedia/Files/Timologio/web_portal_english_timologio_loipwn_ergasiwn.pdf?rev=a4a380ff93f84cbbae69f0cf70a6bf69&hash=A5262D11D5AF9449D45861FE0CDDFF3D6.

¹⁸ <https://www.reuters.com/article/lifestyle/greek-burglars-cash-in-as-savers-flee-banks-idUSBRE84N10D/>.

quickly. Individuals who faced higher switching costs still migrated from deposits to those other alternatives, but more slowly. Slower runs can give policymakers sufficient time to act to prevent undesirable outcomes. The case of March 2023 is a cautionary tale of how low switching costs and ease of convertibility can exacerbate incentives to run, and how bank runs can materialize from one day to the next. A CBDC has the potential to substantially reduce switching costs and withdrawal times compared to redeeming in physical currency. As a result, runs on uninsured or very large deposits at banks or other financial institutions may materialize in a very short time, leaving little scope for policy intervention.

3. Mechanisms

In this section, we analyze the mechanisms through which a CBDC can affect financial stability, focusing on the following three broad topics:

- (i) How would a CBDC affect the run dynamics on financial institutions?
- (ii) How would a CBDC affect the credit provision by banks in times of crisis?
- (iii) How would a CBDC help address financial stability risks from stablecoins?

While investigating these questions, we discuss both the negative and positive implications of a CBDC for financial stability. We discuss separately the positive effects of CBDC on stablecoins, because the digital nature of a CBDC could bring additional benefits for them.

The design choices for a CBDC matter for its properties as a safe-haven asset in times of crisis. As mentioned, we consider a CBDC that could be, in principle, widely transferable, be held by any institutional or retail market participant, and be interest-bearing.¹⁹ Throughout our analysis we focus on intermediated

¹⁹Access restrictions, holding or transaction limits, and the level of a CBDC remuneration are important policy tools—rather than rigid design choices. Such

CBDC that is offered through wallets provided by institutions or segregated accounts at custodial institutions.²⁰

3.1 How Would a CBDC Affect the Run Dynamics on Traditional Financial Institutions?

Financial institutions that fund illiquid assets with liabilities redeemable on demand at par are said to engage in liquidity transformation and are exposed to the risk of runs. To analyze the effect of a CBDC on this type of run risk, we consider a stylized model with the following elements:

- Financial institutions are funded by runnable liabilities.
- The runnable liabilities can have money-like attributes and provide payment services besides their role as a store of value.
- These financial institutions hold both safe/liquid and risky/illiquid assets.
- Concerns about losses on the risky/illiquid assets can result in self-fulfilling runs whereby holders of runnable liabilities withdraw their funds out of fear that others will do the same.

The model description and analytical results we derive can be found in the appendix. The model is a modified version of the bank run model in Goldstein and Pauzner (2005).²¹

tools can mitigate potential effects of a CBDC on financial stability and will be examined in Section 4.

²⁰In our analysis a CBDC is a liability of the central bank that is distributed/allocated/transferred to the public with the help of institutions rather than via direct individual accounts at the central banks.

²¹Ahnert et al. (2023) also study the effect of a CBDC on bank run risk using a similar bank run model. They examine how banks may respond to higher CBDC remuneration by increasing deposit rates, which yields ambiguous effects, as banks are able to mitigate run incentives for low-enough levels of CBDC remuneration. This mechanism is absent in our model because all bank revenues are distributed to deposits, as in Goldstein and Pauzner (2005), and there is no room for the bank to adjust its profits and deposit rate to mitigate run incentives. Indeed, Ahnert et al. (2023) show that the ambiguity stems from banks' market power in the deposit market, and under perfect competition CBDC remuneration unambiguously increases run risk. Complementary to them, our model focuses on the role of liquidity transformation and payment services of deposits on run risk as in Kashyap, Tsomocos, and Vardoulakis (2024), which are elements absent in Ahnert et al. (2023) and which matter, as we show, for the implication

We examine various channels through which a CBDC can affect run risk at financial institutions that engage in liquidity transformation. Note that not all channels are equally important for every institution, and we provide examples of institutions for which each channel is relevant.

3.1.1 Channel 1: A CBDC Provides an Attractive Option for Flight to Safety

The level of run risk at financial institutions depends on the degree of liquidity transformation, as measured by the asset share of liquid assets, and on the availability and desirability of an outside option to run to. For example, large uninsured depositors or investors at prime MMFs may redeem their deposits or shares, respectively, and place them in government MMFs, while small retail depositors could also choose to redeem their bank deposits in currency. A CBDC may offer a more attractive alternative, which could intensify the risk of runs at financial institutions.

The attractiveness of a CBDC as a safe haven accrues not only from the safety inherent in central bank liabilities but also from other important characteristics, such as the payment services offered by a CBDC, the low cost of switching to and from a CBDC, and the interest rate that the central bank may choose to pay on CBDC holdings. For example, large uninsured depositors may find it more convenient to hold CBDC during a crisis compared to transferring their deposits to an alternative bank or investing funds in government MMFs, given that switching costs to CBDC could be smaller and that the CBDC could be used directly for transactions. Both benefits can be considerable. The few available studies on the subject suggest that switching costs for moving one's deposits to another bank are not negligible; a CBDC could effectively reduce such costs.²² Moreover, the incremental benefit of using a CBDC in

of remunerated CBDC on run risk. See also Schilling, Fernández-Villaverde, and Uhlig (2024) for a nominal bank run model to examine the trade-off introduced by CBDC between bank runs and price stability.

²²Shy (2002) uses Finnish banking data with a simple calibration method. The estimated switching costs for moving one's deposits into another bank account for between 0 percent and 11 percent of the average balance a depositor maintains with the bank. Stenbacka and Takalo (2019) also use the same method with

transactions could be approximated by the difference between three-month Treasury bills and deposit rates, which may represent the relative convenience yield of the latter. This benefit depends on the level of rates, and in recent years has been around 80 basis points.²³ On top of these characteristics, an attractive CBDC remuneration can further increase the flight-to-CBDC incentives.

We can group these distinct incremental benefits—reduced switching costs, convenience from use in transactions, interest paid on CBDC—together in an aggregate incremental benefit and use the aforementioned model to show how run incentives increase as the CBDC becomes a more desirable outside option. We consider three stylized financial institutions: (i) an institution with high liquidity transformation, (ii) an institution with low liquidity transformation, and (iii) an institution with low liquidity transformation and liabilities that can be used as means of payment.

There are a couple of interesting takeaways from this analysis, the details of which can be found in the appendix. First, not surprisingly, run risk is increasing in the attractiveness of the CBDC as an outside option. Second, the incremental increase in run risk is higher for institutions with higher liquidity transformation because these firms hold fewer assets that are liquid in all stress events compared to their runnable liabilities. Third, the destabilizing effect of a CBDC is much smaller for institutions with liabilities that are useful as means of payment. Taken together, our analysis suggests that the run risk of bank-like institutions with low liquidity transformation is expected to be small. However, the introduction of a CBDC may increase run risk at fragile arrangements, such as prime MMFs.²⁴

more recent Finnish deposit market data. Their estimated switching costs are approximately 50 percent higher than those of Shy (2002). Kim, Kilger, and Vale (2003) use Norwegian banking industry data and estimate the magnitude and significance of switching costs. The point estimate of the average switching cost is 4.1 percent, about one-third of the market average interest rate on loans.

²³See Van den Heuvel (2019).

²⁴As noted previously, competitive pressures from CBDC could discipline banks, urging them to adapt their business model appropriately (for example, Chiu et al. 2023). We are abstracting away from such a consideration in our analysis that could weaken the first-order effect we highlight. Nevertheless, we also abstract away from factors that could strengthen our results, such as the fact that CBDC could enhance interoperability, amplifying contagion channels in times of stress and exacerbating the run dynamics we highlight.

3.1.2 Channel 2: CBDC Could Affect Institutions' Reliance on Wholesale Funding

A CBDC may alter the composition of the liabilities of financial institutions in normal times, affecting both banks and nonbanks.

The introduction of a CBDC may result in a shift in bank liabilities from stable insured bank deposits to other funding sources which may be less stable, should depositors find it more convenient to hold CBDC.²⁵ Banks are subject to liquidity requirements such as the liquidity coverage ratio (LCR), which broadly requires that a bank's high-quality liquid assets (HQLA) are higher than the appropriately weighted runnable liabilities within a 30-day horizon. An increase in nondeposit liabilities may require additional HQLA—depending on existing buffers over the LCR minimum requirement—which would decrease bank profitability and increase the cost of intermediation. But the effect on banks' run risk may be muted, especially if the composition of HQLA is tilted to short-term liquid instruments such as reserves.²⁶

Nonbanks may find it more expensive to raise secured wholesale funding should investors find that CBDC offers a more attractive safe investment, inducing nonbanks to turn to other, potentially more fragile, sources of wholesale funding. The implications of this shift may not be prohibitively high during normal times. Yet, it can have an adverse effect on financial stability if institutions enter a crisis with a higher share of runnable liabilities.

A contrarian argument suggests that the introduction of a CBDC, especially if remunerated, has the potential to satiate the demand for safe assets and suppress convenience yields of private

²⁵ For example, Whited, Wu, and Xiao (2022) show that banks can replace a large fraction of deposits with wholesale funding, making banks' funding costs more sensitive to changes in short-term rates.

²⁶ We should note that the LCR weights for short-term liabilities were calibrated according to runoff rates absent a CBDC. The weights capture the haircuts applied to various short-term liabilities that enter the calculation of the LCR, determining the amount of HQLA needed to support a dollar of runnable liabilities. Should a CBDC render some short-term liabilities more prone to runs, adjustments to LCR calibration may be needed to neutralize the effect of a compositional change in institutions' liabilities on run risk and, thus, on financial stability.

short-term liabilities. As a result, incentives for financial and non-financial firms' potentially excessive reliance on short-term funding could be eliminated, enhancing financial stability. See Carlson et al. (2016) and Greenwood, Hanson, and Stein (2016) for a non-CBDC-specific exposition of this argument. We should note that the ability of the central bank to satiate the demand for safe assets depends on how it sterilizes the issuance of a CBDC. If CBDC issuance is accompanied by purchases of Treasury bills by the central bank, or by just replacing banks' reserves, then the net effect on the supply of publicly issued safe asset would be small, potentially close to zero.²⁷

3.1.3 Channel 3: CBDC Could Enhance Liquidity Management at Institutions

Although banks and certain nonbanks have access to central bank liabilities, such as reserves or the ON RRP, many other financial institutions do not. These institutions nonetheless may have large and variable dollar payment needs to manage for which they must rely on other institutions. A CBDC could widen access to a risk-free store of value and settlement asset and could enhance the ability of some financial institutions to manage their liquidity. If such widened access included foreign institutions and/or close interlinking with other central banks, CBDC could be a catalyst for cross-border and cross-currency liquidity risk management benefits. This potential benefit, however, does not appear unique to CBDC, but rather is a function of access to digital central bank liabilities and payment services more generally. Similarly, any negative trade-offs attributed to expansion of allowable access would also be equally relevant to a CBDC.

Using the same model described earlier, we can reevaluate the effect of the introduction of a CBDC on run risk, considering that some of the CBDC benefits, such as remuneration and payment services, accrue to financial institutions holding it as a liquid asset in their portfolios. The risk of runs may continue to increase with CBDC benefits, but it could be lower than it would have been if the institution did not hold CBDC, especially for institutions with a low

²⁷This underscores the importance of how the central bank *recycles* CBDC inflows back to the economy. See Infante et al. (2024) for a detailed discussion.

degree of liquidity transformation, because the institution's liquidity position also improves through the benefits offered by CBDC. In certain cases, the benefits offered by a CBDC may even dominate the adverse effect on run incentives, and run risk may be lower compared to the case in which an institution does not derive additional benefits from holding CBDC among its liquid assets.

3.1.4 Channel 4: CBDC Could Provide Policymakers with Valuable Information during a Run

CBDC could provide policymakers with valuable information that can be used to ameliorate undesirable outcomes stemming from a flight to safety. One potential example is related to the ability of the central bank to monitor the flow of funds into CBDC in real time. Setting aside the nontrivial operational challenges, such monitoring could allow policymakers to identify and resolve weak banks sooner, potentially increasing depositor confidence and making depositors less prone to running on the banks.²⁸

3.2 How Would a CBDC Affect Credit Provision by Banks in Crisis?

As discussed earlier, a permanent shift of funds from deposits to CBDC may require banks to rely on riskier or more expensive wholesale funding markets to a greater extent than they currently do. A greater reliance on wholesale funding relative to deposits could increase the cost of funds to banks.²⁹ In turn, banks may increase their lending rates, affecting credit availability across the economy.³⁰ As a result, the introduction of a CBDC might restrict banks' ability to extend credit, resulting in either a reduction in overall credit or in disintermediation of the banking sector.³¹ Such concerns may be

²⁸See Keister and Monnet (2022).

²⁹Wholesale funding that is secured by HQLA is not necessarily a more costly source of funding for banks than deposits, as that depends on their relative convenience yields and the perceived safety of deposits, if uninsured.

³⁰See Keister and Sanches (2022).

³¹The analysis below assumes that banks pass through the increase in funding costs to borrowers in times of stress. However, Chiu et al. (2023) show that if banks have market power in the deposit market, a CBDC can enhance competition, expanding intermediation. We should note that their results may apply

more significant in a stress event when a reduction in bank credit can result in a credit crunch exacerbating the stress.

Banks that experience a deposit loss to CBDC have a set of options to substitute for lost funding. They could resort to the discount window or turn to secured or unsecured wholesale funding markets. Banks could also decide to raise long-term debt or retain more of their earnings to boost their equity, even though doing so might be costly or impossible during times of stress³² Figure 5 shows the incremental pecuniary cost of borrowing from the discount window or wholesale funding markets over short-term safe rates—left and right panel, respectively, over the past two decades. Both costs are not negligible, and we will consider them under different scenarios.³³


Consider a severe stress event that finds banks well capitalized and with adequate liquidity. As already mentioned above, we do not consider cases in which banks are the source of the stress because the resulting drop in credit extension could be then better attributed to insufficient banking capital rather than to a CBDC. Moreover, credit extension may fall because of higher risk premia reflecting the lower growth prospects. What we are interested in is the incremental effect of a deposit flight to CBDC on credit extension. Thus, our estimates

better to normal times rather than times of stress, which is the focus of our paper. We are solely focusing on the changes to financial stability and not taking a stance on whether the equilibrium allocation will be more efficient, as Chiu et al. (2023) show, or the other way around, as Whited, Wu, and Xiao (2022) show. Even if a CBDC can enhance allocative efficiency in the banking system, it may not be the best solution to deal with the issue of allocative inefficiency.

³²However, banks may not have earnings or be in a position to raise fresh equity, as was the case during the 2007–09 financial crisis. In such circumstances, banks may need to deleverage, which significantly reduces credit supply. We do not examine this case, because the cause of deleveraging would arguably not be a CBDC but rather solvency concerns at banks, which is an important but separate issue.

³³Central banks could decide to set the rate on their lending facilities equal to prevailing deposit rates and allow private banks to replenish the shortfall of deposits with central bank lending, leaving private bank lending rates unaffected. This type of policy is similar to the equivalence result of Brunnermeier and Niepelt (2019), which characterizes how banks can borrow from the central bank to facilitate a swap between private and public money, while leaving prices and equilibrium allocations unchanged. However, there may be significant institutional restrictions—inspired, in part, by the Bagehot principle—that may limit central banks' ability to lend freely at prevailing deposit rates.

Figure 5. Borrowing Costs in Stress Times



should not be regarded as capturing the total effect on credit extension during times of stress, but rather the incremental effect due to the introduction of a CBDC. We consider three possible scenarios.

- (i) In the first scenario, banks borrow from the discount window at an extra cost of about 50 basis points, which is the conditional mean during crisis periods in the last two decades.
- (ii) In the second scenario, banks make up for their lost deposits with various types of wholesale funding, with an average additional cost of about 100 basis points.³⁴
- (iii) In the third scenario, banks replace lost deposit funding with unsecured wholesale funding, reaching an additional cost of funding of about 250 basis points, which is the upper end of the historical realization during the 2007–09 financial crisis.

Table 1 reports the incremental effect of banks' increased cost of funding on the volume of bank lending due to a flight to CBDC in times of stress. Our focus is on the effect of a marginal loss of deposits to CBDC, which banks substitute with costlier funding. That is, if banks had to make up for a dollar of deposits directed to lending in times of stress with more expensive funding source, then they would borrow less and extend less credit to be able to charge a higher lending rate, compensating for the higher funding costs. Implicitly, we assume that banks would extend new credit in times of stress first using deposits, then using the discount window or wholesale funding. These cases span our scenarios. The situation might be different in the steady state whereby some lending would be supported by equity; thus, the effect of a marginal loss of deposits on credit supply would have to be weighted by their share in total funding sources.

We assume that banks would pass through the extra cost to borrowers rather than absorbing part or all of it, which would directly reduce their profitability. As such, the estimates for the decline in

³⁴This estimate is based on Bassett and Rappoport (2022), who compute a comprehensive cost of short-term wholesale funding—as defined in the U.S. implementation of the capital surcharge for globally systemically important banks—which includes (i) secured funding transactions, (ii) unsecured wholesale funding, (iii) covered asset exchanges, (iv) short positions, and (v) brokered deposits.

Table 1. Decrease in Lending during Stress Times

Type of Credit	% Change in Lending for Deposits' Flight to CBDC		
	Scenario (i)	Scenario (ii)	Scenario (iii)
Mortgage Credit Semi-elasticity of Demand about 2.5 (De Fusco and Paciorek 2017)	-1.25%	-2.50%	-6.25%
C&I Loans Semi-elasticity of Demand about 2 (Bassett et al. 2014)	-1.00%	-2.00%	-5.00%
Consumer Credit Semi-elasticity of Demand about 5.2 (Gross and Souleles 2002)	-2.60%	-5.20%	-13.00%

credit provision reported in Table 1 could be regarded as a lower bound.

We examine the effects on different types of loans using estimates for the semi-elasticity of loan demand to interest rates from the literature.³⁵ Scenario (ii) or a scenario between (ii) and (iii) are arguably the most likely ones, implying, for example, a decline in C&I loan provision between 2 and 5 percent in times of stress.³⁶

Well-capitalized banks can be a source of stability by providing liquidity to struggling firms through the provision of lines of credit.³⁷

³⁵The semi-elasticity of loan demand to interest rates captures the percentage decline in loan demand for a 1 percentage point increase in interest rates; these elasticities are not exclusively derived for crisis periods.

³⁶The effect of deposit substitution on credit provision should be considerably lower in normal times. Considering an average TED spread of 30 basis points during normal times—spanning from 2010 to 2019—and a maximum deposits runoff rate of 40 percent, we obtain an increase in the average cost of funds of 12 basis points, which implies a decline of 0.30, 0.24, and 0.62 percent in mortgage credit, C&I loans, and consumer credit, respectively.

³⁷Li, Strahan, and Zhang (2020).

However, adverse effects on their willingness to provide such credit during a stress event, owing to higher funding costs from introducing a CBDC, could amplify a stress event. The combination of a potential flight to CBDC by depositors with a higher reliance on riskier wholesale funding may render the provision of credit lines unprofitable *ex ante*, resulting in lower liquidity insurance and instability during a crisis.³⁸

3.3 How Would a CBDC Affect Financial Stability Risks from Stablecoins?

Despite the original introduction as alternatives to national currencies, cryptocurrencies like Bitcoin have exhibited extreme price volatility measured in dollars. This volatility has made them less appealing as a medium of exchange or store of value. Stablecoins emerged to serve as a relatively stable store of value within the crypto space, offer lower-cost entry and exit points to crypto markets, and are the main collateral asset for leveraged positions in the digital asset ecosystem.³⁹ Total issuance of stablecoin continues to grow. However, stablecoins have experienced runs, either because of their structure or because of doubts about their backing.⁴⁰

A widely accessible CBDC might help to mitigate the instability introduced by stablecoins, as it could supplant stablecoins that are relatively risky or improve the quality of some stablecoins' reserve assets.⁴¹ If stablecoins were backed one-to-one by a CBDC, they would not be as susceptible to runs, though CBDC would not mitigate issues arising from legal, custodial, or operational concerns.⁴² Stablecoin issuers might then compete on technical innovation rather than on relative perceived capacity to redeem. The introduction of a CBDC to which stablecoin issuers had large-scale access thus could encourage the development of private stablecoins. Potentially, stablecoins arrangements may prefer CBDCs to other alternatives such as

³⁸See, also, Piazzesi and Schneider (2022).

³⁹Gorton et al. (2026).

⁴⁰See Azar et al. (2022) for a review of financial stability risks in the digital asset ecosystem, including stablecoins.

⁴¹See, among others, Gorton and Zhang (2023).

⁴²Consistent with earlier discussion of firm liquidity risk management, access to a CBDC granting a master account at the Federal Reserve or giving access to the ON RRP could all have similar implications for stablecoin stability.

Treasury bills or other central bank liabilities as safe reserve assets. If the design of a CBDC offered greater interoperability within the digital asset ecosystem than these alternatives, it could reduce transaction costs when moving between digital and traditional financial platforms and allow for more efficient settlement of digital transactions. In addition, depending on design choices related to information management, CBDC transaction records could be made publicly observable in real time, improving transparency of those stablecoins' reserves.⁴³

4. Policy Tools to Address Financial Stability Concerns of CBDC

A variety of tools may be available to address the financial stability concerns associated with a potential disruptive flight to CBDC in times of stress. All the tools proposed by academics and policymakers alike attempt to limit the desirability of CBDC by reducing its appeal or availability as a store of value. We focus on two broad categories: (i) price tools, under which the central bank controls the remuneration on a CBDC, and (ii) quantity tools of various types, including tools limiting the balances of CBDC that holders are allowed to maintain in their wallets/accounts.

Remuneration and programmability are two design choices that can facilitate the use to tools to mitigate financial stability risks for a CBDC. Designing a CBDC to be remunerated enables the adoption of price tools to discourage a flight to safety in times of stress. Designing a CBDC to be programmable may provide built-in rules that constrain its use.⁴⁴ A programmable CBDC potentially can overcome implementation hurdles associated with certain quantity tools as well as enable a novel category of tools. Table 2 summarizes the set of tools we examine, highlighting the effect of each tool on financial stability as well as possible undesirable side effects.

⁴³See Chiu and Monnet (2024) for a rigorous analysis of the interaction between CBDC and stablecoins, together with tokenized bank deposits, with respect to customer protection, financial stability, and illicit activities.

⁴⁴Two natural components of the definition of programmable money are a digital form of money and a mechanism for specifying the automated behavior of that money through a computer program. See Lee (2021) for more details on programmable money.

Table 2. Tools to Mitigate Flight to CBDC

Tool	Type of Tool	Effect on Run Incentives	Potential Adverse Side Effects
Tiered Remuneration	Price Tool	Discourage flight to safety for moderate stress levels	Discourage use of CBDC as a medium of exchange during a crisis if not implemented properly
Access Limits	Quantity Tool	Reduce risk of runs originating from the entities on which the restrictions are placed	Harm the medium of exchange role of CBDC, increase the incentives for regulatory arbitrage
Transaction Size Limits	Quantity Tool	Discourage flight to safety for moderate stress levels	Discourage use of CBDC as a medium of exchange if not implemented properly
Holding Limits	Quantity Tool	Prevent hoarding of CBDC above the limits in all instances	Ineffective if the limit is not chosen adequately for different entities; regulatory arbitrage; amplification of runs if set at aggregate level
Convertibility Limits	Quantity Tool	Stop a run if their implementation is unanticipated	Preemptive runs; multiple secondary markets for CBDC

We study the ability of each potential tool to mitigate harmful runs on financial institutions during times of crisis. We highlight the possibility of calibrating these tools differently during normal times and crisis times to avoid restricting CBDC demand during normal times. No tool is individually superior to others, and each tool introduces trade-offs with respect to its costs, benefits, and ease of implementation. Finally, we discuss other tools that central banks could deploy to safeguard financial stability from a flight to CBDC.

4.1 Price Tools Relevant for Interest-Bearing CBDC

The first category of tools addresses the financial stability concerns of a flight to CBDC by directly targeting the interest rate of a CBDC. We refer to such tools as price tools, which can be made state-contingent to more efficiently target flight to CBDC concerns in times of stress, while preserving the attractiveness of CBDC in normal times.

The simplest form of such tools is a single-tiered remuneration scheme, where the interest rate on CBDC holdings is the same irrespective of the amount of CBDC. According to this scheme, the interest rate on CBDC holdings needs to be sufficiently low to make a CBDC relatively unappealing and to discourage a potential flight to safety in times of stress. Importantly, a CBDC would provide transaction services and other benefits, in addition to interest payments. As a result, dissuading a run to a CBDC may require a negative interest rate, potentially very low in some cases, highlighting how a CBDC design that did not include interest payments—that is, a 0 percent interest rate on CBDC holdings—may not be sufficient to limit run risks. Significant or rapid rate changes, particularly into negative territory, may also discourage use of a CBDC as a medium of exchange, which may further deter a flight to CBDC but may worsen welfare if CBDC payments are valuable in times of stress.

A more elaborate price tool is a two-tiered remuneration scheme, where CBDC holdings below a threshold, or in the first tier, are remunerated at a higher interest rate than those above the threshold, or in the second tier.⁴⁵ The interest rate differential between

⁴⁵See Bindseil, Panetta, and Terol (2021). Note that this remuneration scheme could be designed with multiple tiers.

larger and smaller CBDC holdings aims to discourage large holdings of CBDC as a store of value in times of market stress while minimizing effects on its role as medium of exchange. In times of severe stress, sizable negative rates may be required to prevent a flight to CBDC, regardless of tiering.⁴⁶ For moderate stress scenarios, however, two-tiered remuneration might reduce the likelihood of a flight to CBDC. This type of tiered remuneration would mimic the price elasticity of the demand for other safe-haven assets. Hence, all else equal, the likelihood of a run on financial institutions induced by a flight to CBDC is lower.

When implementing a two-tiered remuneration scheme for a CBDC, selecting the threshold separating the first from the second tier might be challenging. In fact, the choice of threshold by the central bank requires significant investment in gathering information about the interest rate elasticities of the different entities and individuals holding CBDC. Institutional investors and retail depositors may have different needs for transactions; thus a universal threshold might be inadequate to deter runs for both types of investors while allowing them to use CBDC for transactions. Moreover, the potential financial stability benefits of CBDC described above—for example, those related to improving the safety of the stablecoin industry—would be constrained by the imposition of such universal threshold. In these cases, a version of a tiered remuneration scheme based on voluntary remuneration targets (VRTs) can provide a solution.⁴⁷

⁴⁶It is worthwhile to note that for depositors who can hold cash, a natural constraint on the remuneration rate for the second tier is given by the remuneration rate of their outside option, such as cash. Thus, lowering the remuneration rate on the second tier would be ineffective at discouraging the run if depositors want to run to cash and there are no other limits on their ability to do so. The introduction of a CBDC, however, is not relevant for this type of run, which could occur even in absence of a CBDC. Moreover, note that cash may only be a viable outside option for low-enough deposit balances.

⁴⁷A VRT is a form of tiered remuneration where CBDC holders communicate to the central bank a target for their desired CBDC holdings and must pay a penalty should they miss it. Because holdings below the target are remunerated at the higher rate on the first tier while those above the target are remunerated at the lower rate on the second tier, CBDC users have an incentive to choose a nonzero target. At the same time, the penalty for missing the target incentivizes users to set a target equal to their expected holdings. See Baughman and Carapella (2020, 2023).

Finally, it should be noted that even the simplest form of CBDC remuneration poses significant challenges for policy implementation. Accurate estimates of the elasticity of CBDC demand to its remuneration will depend on its uses, which, in turn, are intrinsically linked to its design features and to the evolution of the broader payment landscape.

4.2 Quantity Tools Relevant for Interest-Bearing and Noninterest-Bearing CBDCs

The second category of tools to address a flight to CBDC place scale restrictions on activities that users can undertake with a CBDC. We refer to such tools as quantity tools.

4.2.1 Access Limits

Who may access CBDC is a key question that likely would be determined by the intended purpose of the CBDC along with other legal considerations not examined in this paper. Access policy choices would also directly shape who has an option to run into CBDC in stress. To the extent that withdrawals by specific entities are more likely to trigger a run on banks or other financial institutions, prohibiting those entities from holding CBDC would reduce the risk of runs. For example, prohibiting institutional investors from holding CBDC would result in a reduced risk of banks and MMFs experiencing runs from large institutional deposits/investment.

4.2.2 Transaction Size Limits

Transaction size limits would constrain how much CBDC can be sent in a single payment. Such limits are common features of payment systems. In principle, there could be multiple, differentiated CBDC payment systems, as is the case for balances in Reserve Bank accounts, where transactions occur using multiple small- and large-value payment systems depending on the needs of the sending and receiving banks.⁴⁸ As in existing payment systems, a key

⁴⁸For example, the Fedwire Funds Service has a transaction size limit of just below \$10 billion and the default transaction limit for the FedNow Service is \$100,000 with the option to adjust up or down with a current ceiling of \$500,000. Both Fedwire and FedNow are real-time gross settlement payment systems, but

rationale for potentially limiting the size of individual transactions using CBDC is to prevent or detect and to lessen the consequences of accidental or fraudulent transfers, or of cyberattacks.⁴⁹ An additional perspective on transaction size limits is that they can be in practice equivalent to access limits, yet less distortionary. If desired, reasonably small transaction size limits could allow individuals and commercial entities to use CBDC efficiently for a variety of consumer and commercial purposes while substantially reducing CBDC's usefulness for financial institutions' broader financial-sector activities. Hence, financial institutions might be discouraged from holding large quantities of CBDC due to the inconvenience of using it in large amounts. Such a transaction size limit would be less distortionary than an access limit, as it would still be possible for financial institutions to hold, send, and receive CBDC making smaller payments: use would not be prevented, but simply impractical for large, financial market purposes. Nonetheless, transaction size limits would not affect the usefulness of CBDC as a store of value in times of stress and, thus, would likely not be useful to deter widespread runs in times of severe stress.

4.2.3 Holding Limits

Holding limits would restrict the amount of CBDC that individuals or entities can hold or trade.⁵⁰ Similar to ON RRP caps discussed earlier, holding limits could be individual or aggregate.

Aggregate limits would cap the stock of CBDC outstanding and could control the aggregate outflows from the financial system during times of crisis, thus mitigating the financial stability concerns from runs on financial institutions. They may, however, create incentives for preemptive runs, as investors may fear that the aggregate limit is reached quickly and that they will not be able to fly to CBDC once the crisis is imminent. For example, an individual depositor

they differ in transaction size limits and other design features to better serve different primary use cases and their respective risk management needs. Similarly, the Reserve Banks also own and operate the National Settlement Service and ACH payments systems, which offer alternative settlement methods to meet user needs.

⁴⁹For example, see Duffie and Younger (2019) and Rubinfeld (2019).

⁵⁰See, for example, Bindseil (2020).

may want to withdraw their deposits due to fears that the aggregate limit might be met through redemptions of other depositors, leaving them unable to obtain CBDC unless they withdraw first. By contrast, individual limits, placed at the level of wallet or physical individual/entity, could address run incentives more directly than aggregate limits, as they could be targeted to the typical size of activities of a given wallet/entity. Moreover, individual limits do not link individual run incentives to the probability that others might withdraw CBDC in sufficiently large amounts to meet the aggregate limit. Hence, individual limits would not induce incentives for coordination.⁵¹

One approach to limit the desirability of CBDC as a safe haven without damaging its role in facilitating transactions would be to carefully calibrate individual holding limits. Such calibration poses challenges similar to the choice of remuneration thresholds for tiered remuneration. Given uncertain demand for CBDC beyond its role as a safe haven, setting holding limits may involve a process of trial and error. For example, the choice of holding limits for financial institutions could be informed by the experience of setting individual limits for other central bank facilities available to nonbanks, such as the ON RRP facility. Similarly, the choice of holding limits for retail depositors could be set below the deposit insurance limit for bank deposits to minimize the role of CBDC as a safe haven.

We should note that quantity limits could be combined with tiered remuneration to further improve outcomes in times of stress. As an example, suppose that the individual holding limit was sufficiently large and above the remuneration threshold separating the two remunerated tiers. CBDC holdings in the higher-remunerated tier could be unremunerated, while those in the lower-remunerated tier could be charged a penalty rate and those above the holding limit would not be allowed. This approach would allow users

⁵¹See Huang and Keister (2024) for an analysis of preemptive runs on MMFs and its application to the outflows experienced by MMFs in March 2020. Consistently with our argument, Huang and Keister (2024) find that investors may have an incentive to run preemptively when redemption fees on MMF shares are imposed only once redemption demand is unusually high. To be effective in preventing runs, the fee could apply even when redemptions are in the normal range, despite the cost in terms of welfare from the reduction in the liquidity-provision role of the fund.

who would like to transact with CBDC to do so up to the holding limit, though at a cost when balances rose to the level of the lower-remunerated tier. Such costs would discourage transfers into CBDC above the lower-remunerated threshold and deter widespread runs, while allowing the CBDC to facilitate payments.⁵²

4.2.4 Convertibility Limits

Convertibility limits at the account holder level are restrictions that could be imposed on flows out of financial institutions and into CBDCs. Most simply, they would be constructed as a velocity limit, constraining how much can be converted from one type of money to another in a given period of time. A cap on cash withdrawals from ATMs is an example of an existing convertibility limit. The potential drawback of convertibility limits is that they may lead to preemptive, slowly unfolding runs as investors may attempt to move funds out of their financial institutions within their convertibility limit. Hence, while convertibility limits might be effective at slowing down the pace of sudden, unanticipated flight-to-CBDC episodes, they may amplify runs triggered by noisy signals about the limit being met, thus contributing to a more prolonged surge in requests of withdrawals.

4.3 Implementation Hurdles and CBDC Programmability

Most of the aforementioned tools, especially quantity limits, share concerns common to the effective implementation of monetary and fiscal policies, but also raise some novel challenges.

The first is an inherent time-inconsistency problem of all policies that require commitment by policymakers to carry them through. For example, the central bank may have an incentive to increase CBDC quantity limits in the event of a developing run, in order to

⁵²Effectively, the higher-remunerated tier would be equivalent to an exemption threshold from negative remuneration rates on CBDC. This approach leverages the idea that quantity limits are economically equivalent to tiered remuneration, as the threshold separating the first remuneration tier from the second identifies the quantity of allowed balances held in an individual account or wallet, thus remunerated at a possibly zero net interest rate, while balances above the threshold are remunerated at a -100 percent rate, thus they are not allowed.

avoid damaging the transaction role of CBDC or disfavoring depositors without alternate options to store their wealth safely.⁵³ If the central bank lacks commitment or credibility to implement certain tools when opposing incentives are present, then investors may anticipate the policy change and their motives to run will be less affected by the quantity limit.⁵⁴

Second, a potential drawback of some quantity tools is encouraging the creation of secondary markets for CBDC, which might result in different prices quoted for CBDC and other central bank liabilities. If the central bank cannot ensure one-to-one convertibility of CBDC into other forms of central bank money, the parity among all forms of central bank money would depend on arbitrage, potentially leading to price dislocations in times of stress. Consider the case of two market participants who would like to hold CBDC but one of them has already reached the holding limit. They could bypass the holding limit by creating another digital asset—for example, a wrapped CBDC.⁵⁵ A wrapped CBDC would be a new token, issued by a smart contract that retains custody of the actual CBDC until the wrapped CBDC is returned. Wrapped CBDC can then be traded in decentralized exchanges or in liquidity pools, where its role as a store of value would be as good as that of actual CBDC.⁵⁶ The reason is as follows. If wrapped CBDC was traded below the price at which actual CBDC is traded, then any market participant—not facing a binding holding limit—could purchase it, then return it to the smart contract to unlock the stored CBDC, thus realizing a profit equal to the difference in prices between the wrapped and actual

⁵³In this case the policy or the tool is defined to be time inconsistent (Kydland and Prescott 1977).

⁵⁴For an application to suspension of convertibility to prevent bank runs, see Ennis and Keister (2009).

⁵⁵Wrapped tokens are a way to use cryptocurrencies on blockchains other than the blockchain on which they were originally built, thus proposing a solution to the problem of blockchain interoperability. Wrapped tokens are backed one-to-one by their underlying asset, which is stored in a digital vault.

⁵⁶In this scenario, we implicitly assume that wrapped CBDC could be minted in distributed ledgers using smart contracts, as wrapped cryptocurrencies are currently minted, and that CBDC also could be traded on distributed ledgers or against tokenized assets that can themselves be traded on distributed ledgers and held in smart contracts. The central bank may find these arrangements desirable to enhance interoperability or to provide an alternative to private stablecoins.

CBDCs. Note that the actual CBDC is never owned by the market participant with a binding quantity limit. Rather, this market participant will hold a wrapped CBDC and hence there is no violation of individual quantity limits in terms of actual CBDC.⁵⁷

More generally, all quantity tools could be potentially circumvented by designing securities to be issued and traded outside the purview of policymakers to mimic the forbidden CBDC trade. The issuance and trade of such securities may not be prevented unless policymakers have detailed information about the public's transaction behaviors and payment flows.⁵⁸ As a result, a secondary market for CBDC may arise, which might lessen the effectiveness of the quantity tools. Moreover, tools may be difficult to implement adequately. This is particularly true for quantity tools, as they would require a lot of information about the identity and the activity of CBDC users.

A CBDC designed with programmable features may introduce ways to potentially overcome some of these hurdles. For example, the problem of time inconsistency can be addressed by programming individual limits in CBDC together with costs embedded in the CBDC. While the introduction of such costs would not directly force commitment onto the central bank and policymakers, it would reduce the severity of their time-inconsistency problem by reducing the benefits of bypassing the limits.⁵⁹

⁵⁷Note that the degree to which counterparty or other additional risks are entailed in this transaction will depend on its specific implementation.

⁵⁸It is worthwhile to highlight that the solution to the problem of regulatory arbitrage requires the tools developed by the literature on (dynamic) mechanism design with hidden trades/actions. Despite the fact that this class of problems are hard to solve, policymakers should be aware that proper design of quantity tools should account for the possibility of hidden actions on the side of the agents.

⁵⁹The literature has suggested additional features for CBDC programmability such as being eligible only for certain transactions during times of crisis—for example, purchases of groceries, medical bills, and similar goods or services—which would effectively make the CBDC useless to be transferred anywhere else and, therefore, to be hoarded. While not citing CBDC explicitly, see Andolfatto (2020) for an application to market freezes. Notice that this notion of programmability could be enabled within several layers of the CBDC architecture, such as the platform, coin/token, contract, and wallet level, all of which would be characterized by different trade-offs. See Lee (2021).

4.4 *Ex Post Tools to Preserve Financial Stability*

The tools described so far can be characterized as *ex ante* tools aimed to promote financial stability by discouraging activities that can endanger it ahead of any stress event. It is also worthwhile to highlight a wide set of *ex post* tools aimed to preserve financial stability once stress hits. Rather than being specific to the flight-to-safety concern arising from the introduction of CBDC, such tools belong to policymakers' machinery to respond to any financial crisis regardless of the event that triggers it. Central bank lending is an example of such a tool; in the event of a run that causes banks to be illiquid, access to the central bank lending provides ready access to funding and can therefore help avoid potential insolvency of an individual bank. In general, all the tools traditionally adopted by central banks for crisis response would be available after the introduction of a CBDC. The economic mechanisms by which such tools would operate are similar to the mechanisms by which they act in a world without a CBDC. The scope of such tools, however, might be broader, as the introduction of a CBDC could bring along a new set of assets and operations that were not technologically feasible before.

Appendix

We consider a simple bank run model akin to Goldstein and Pauzner (2005) extended to account for both liquid and illiquid asset holdings and payment services offered by deposits similar to Kashyap, Tsomocos, and Vardoulakis (2024).

Consider an economy with three time periods, $t = 0, 1, 2$, one bank, and a continuum of representative, risk-neutral, depositors of mass one. There is no time discount. The bank is fully funded with deposits D , normalized to one, while competition and free entry drive bank profits to zero such that all bank revenues are distributed to depositors similar to Goldstein-Pauzner. The bank invests the deposits in a portfolio of risky and safe assets. The risky asset requires one unit of investment at $t = 0$ and yields a payoff $R > 1$ with probability $\theta \sim U[0, 1]$ at $t = 2$, and zero otherwise. Following Goldstein-Pauzner, the risky asset may also be illiquid, such that it yields $\xi < 1$ if liquidated early at $t = 1$ for realization of $\theta < \bar{\theta} < 1$,

while it yields R if liquidated early for realization of $\theta \geq \bar{\theta}$.⁶⁰ The safe asset also requires an investment of one at $t = 0$, but it is both safe and liquid, such that it delivers a payoff of one at $t = 2$ and can also be liquidated for one at $t = 1$. Denote by ℓ the share of the safe asset in the bank's asset portfolio, which determines the degree of liquidity transformation. We will not derive ℓ endogenously but rather perform comparative statics on it to examine how run risk from a CBDC differs across banks with different levels of liquidity transformation.⁶¹ Deposits carry a convenience yield from payment services at $t = 2$ equal to V , which depositors only reap if the bank survives.

Depositors receive at $t = 1$ a private noisy signal about the realization of θ given by $x_i = \theta + \epsilon_i$, with $\epsilon_i \stackrel{\text{iid}}{\sim} U[-\epsilon, \epsilon]$. Given the private signal, each individual depositor decides whether to keep or withdraw their deposits in a manner that we detail below. If depositors withdraw early, they forfeit any claim on future bank revenues and just receive the face value of their deposits. If the individual depositor manages to withdraw their deposits successfully, then they can invest the proceeds in a CBDC that promises to repay $\bar{R} \geq 1$ with certainty at $t = 2$ plus a convenience yield equal to \bar{V} . If they keep their deposits at the bank, they receive the total payoff from bank investments distributed pro rata among remaining depositors, plus the convenience yield only when the bank is solvent.

Denote by $\lambda \in [0, 1]$ the portion of depositors that withdraw early. Depending on the level of λ , the number of illiquid assets the bank may need to liquidate could drive it to insolvency at $t = 2$ or may even result in the bank running out of liquidity to serve all early withdrawals at $t = 1$. We derive these two thresholds for λ below. For the remainder of the analysis, we focus on the case $\theta < \bar{\theta}$.⁶²

⁶⁰ $\bar{\theta}$ is an exogenously assumed threshold that defines the region of fundamentals ($\theta \geq \bar{\theta}$) where individual depositors never choose to withdraw irrespective of their beliefs about others' action. For $\theta < \bar{\theta}$ individual actions may depend on beliefs about the actions of others.

⁶¹It would be interesting to introduce bank heterogeneity on different dimensions that could result in different equilibrium ℓ and study how the introduction of CBDC changes these equilibrium levels. We leave this for future work, and focus on the partial equilibrium effects of a CBDC.

⁶²As mentioned, if the payoff to the risky technology is R , which happens when $\theta > \bar{\theta}$, no depositor withdraws and $\lambda = 0$ in equilibrium.

The available liquidity that the bank has at $t = 1$ after serving withdrawals is given by

$$L(\lambda) = \xi(1 - \ell) + \ell - \lambda, \tag{A.1}$$

which implies that the bank runs out of liquidity for $\lambda > \bar{\lambda} = \xi(1 - \ell) + \ell$, in which case withdrawing depositors receive their deposits only with some probability according to sequential servicing.

For $\lambda \leq \bar{\lambda}$ the bank has enough liquidity to serve early withdrawals, but doing so may result in lower revenues at $t = 2$. Given that the illiquid assets are liquidated at a discount, we assume that the bank will first use the liquid assets to serve early withdrawals. But after $\lambda > \ell$ it will also liquidate a portion of the illiquid assets equal to $y = [\lambda - \ell]^+ / [(1 - \ell)\xi] \leq 1$. The depositors that choose not to withdraw receive pro rata the bank revenues, but only enjoy the convenience yields if their deposits are repaid in full (recall that all bank revenues accrue to depositors). With probability $1 - \theta$, the remaining illiquid assets are worth zero at $t = 2$, and depositors do not enjoy a convenience yield. With probability θ , remaining depositors enjoy the convenience yields only if the pro rata distributed payoff they receive is higher than one, i.e.,

$$\frac{R(1 - \ell)(1 - y) + [\ell - \lambda]^+}{1 - \lambda} \geq 1, \tag{A.2}$$

or equivalently if $\lambda \leq \hat{\lambda}$, where

$$\hat{\lambda} = \frac{\xi R(1 - \ell) - \xi + \ell R}{R - \xi}. \tag{A.3}$$

Then, given θ , the payoff differential for an individual depositor between not withdrawing and withdrawing at $t = 1$ to invest in CBDC is given by

$$\nu(\theta, \lambda) = \begin{cases} \theta \frac{R(1-\ell)(1-y)}{1-\lambda} + \frac{[\ell-\lambda]^+}{1-\lambda} + \theta V - (\bar{R} + \bar{V}) & \text{if } 0 \leq \lambda \leq \hat{\lambda} \\ \theta \frac{R(1-\ell)(1-y)}{1-\lambda} + \frac{[\ell-\lambda]^+}{1-\lambda} - (\bar{R} + \bar{V}) & \text{if } \hat{\lambda} < \lambda \leq \bar{\lambda} \\ -\frac{\ell+(1-\ell)\xi}{\lambda}(\bar{R} + \bar{V}) & \text{if } \bar{\lambda} < \lambda \leq 1 \end{cases} . \tag{A.4}$$

Define by $\underline{\theta} = (\bar{R} + \bar{V} - \ell) / (R(1 - \ell) + V)$ the solution to $\nu(\underline{\theta}, 0) = 0$ such that an individual depositor withdraws, i.e., $\nu(\theta, 0) < 0$, for $\theta < \underline{\theta}$ even if no other depositors withdrew.

Given the private signal, an individual depositor will update their posterior about θ , which will be uniform in $[x_i - \epsilon, x_i + \epsilon]$ and compute the expected payoff differential

$$\Delta(x_i) = \int_{x_i - \epsilon}^{x_i + \epsilon} \nu(\theta, \lambda) \frac{d\theta}{2\epsilon}. \tag{A.5}$$

If $x_i \geq \bar{\theta} + \epsilon$, the individual depositor can conclude that $\theta \geq \bar{\theta}$ and will not withdraw, independent of their belief about λ ($\Delta(x_i) > 0$). Similarly, if $x_i < \underline{\theta} - \epsilon$, the individual depositor can conclude that $\theta < \underline{\theta}$ and will withdraw, independent of their belief about λ ($\Delta(x_i) < 0$). These are the *upper and lower dominance* regions for θ , where the individual action is independent of the beliefs about the actions of others.

For intermediate $x_i \in [\underline{\theta} - \epsilon, \bar{\theta} + \epsilon)$, the sign of $\Delta(x_i)$ depends on the beliefs about λ . To pin down these beliefs, we focus on a threshold strategy that all depositors follow. We show that there exists a unique signal threshold x^* , such that every investor withdraws if their private signal $x_i < x^*$ and does not withdraw if $x_i > x^*$. Given this threshold, an individual depositor can form well-defined beliefs about the total number of withdrawals, denoted by $\lambda^b(\theta, x^*)$, and given by the probability that other depositors receive a private signal below x^* . If $\theta > x^* + \epsilon$, all depositors get signals $x_i > x^*$, none withdraw, and $\lambda^b(\theta, x^*) = 0$. If $\theta < x^* - \epsilon$, all depositors get signals $x_i < x^*$, all withdraw, and $\lambda^b(\theta, x^*) = 1$. If $x^* - \epsilon \leq \theta \leq x^* + \epsilon$, some depositors get signals $x_i > x^*$, while others get signals $x_i < x^*$; thus, under the threshold strategy, $\lambda^b(\theta, x^*) = \Pr(x_i < x^*) = (x^* - \theta + \epsilon) / (2\epsilon)$. The following equation summarizes these beliefs:

$$\lambda^b(\theta, x^*) = \begin{cases} 1 & \text{if } \theta < x^* - \epsilon \\ (x^* - \theta + \epsilon) / (2\epsilon) & \text{if } x^* - \epsilon \leq \theta \leq x^* + \epsilon \\ 0 & \text{if } \theta > x^* + \epsilon \end{cases} \tag{A.6}$$

Using (A.6), an individual depositor can compute the expected payoff differential using their posterior about θ , given her signal x_i and an assumed value for x^* :

$$\Delta(x_i, x^*) = \int_{x_i - \epsilon}^{x_i + \epsilon} \nu(\theta, \lambda^b(\theta, x^*)) \frac{d\theta}{2\epsilon}. \tag{A.7}$$

Unlike in (A.5), beliefs in (A.7) are uniquely determined pinning down $\Delta(x_i, x^*)$.

A depositor does not withdraw ($\Delta(x_i, x^*) > 0$) if $x_i > x^*$ and withdraws ($\Delta(x_i, x^*) < 0$) if $x_i < x^*$. By continuity, the depositor that receives the threshold signal x^* is indifferent between not withdrawing and withdrawing, i.e.,

$$\Delta(x^*, x^*) = \int_{x^* - \epsilon}^{x^* + \epsilon} \nu(\theta, \lambda^b(\theta, x^*)) \frac{d\theta}{2\epsilon} = 0. \tag{A.8}$$

A threshold strategy implies thresholds for fundamentals $\theta_{\hat{\lambda}}$ and $\theta_{\bar{\lambda}}$ such that the depositors enjoy the convenience yield at $t = 2$ for $\theta \geq \theta_{\hat{\lambda}}$ and the bank has enough liquidity at $t = 1$ for $\theta \geq \theta_{\bar{\lambda}}$ given signal threshold x^* and withdrawals $\lambda^b(\theta, x^*)$. These thresholds are determined by $\hat{\lambda} = \lambda^b(\theta_{\hat{\lambda}}, x^*)$ and $\bar{\lambda} = \lambda^b(\theta_{\bar{\lambda}}, x^*)$. Using these, (A.8) can be expanded to

$$\begin{aligned} \Delta(x^*, x^*) = & - \int_{x^* - \epsilon}^{\theta_{\bar{\lambda}}} \frac{\ell + (1 - \ell)\xi}{\lambda^b(\theta, x^*)} (\bar{R} + \bar{V}) \frac{d\theta}{2\epsilon} \\ & + \int_{\theta_{\bar{\lambda}}}^{x^* + \epsilon} \left[\theta \frac{R(1 - \ell)(1 - y)}{1 - \lambda^b(\theta, x^*)} + \frac{[\ell - \lambda^b(\theta, x^*)]^+}{1 - \lambda^b(\theta, x^*)} \right. \\ & \left. - (\bar{R} + \bar{V}) \right] \frac{d\theta}{2\epsilon} + \int_{\theta_{\bar{\lambda}}}^{x^* + \epsilon} \theta V \frac{d\theta}{2\epsilon} = 0. \tag{A.9} \end{aligned}$$

As is typical in the global games literature, we focus on the limiting case where noise $\epsilon \rightarrow 0$, which allows us to derive a common threshold for fundamentals, θ^* , such that runs occur for realizations $\theta < \theta^*$. Expressing (A.9) in terms of θ^* and changing variables from

θ to λ , such that as θ decreases from $x^* + \epsilon$ to $x^* - \epsilon$, λ uniformly increases from 0 to 1, we get⁶³

$$\begin{aligned} \bar{\Delta}^* &= \int_0^{\hat{\lambda}} \theta^* V d\lambda + \int_0^{\bar{\lambda}} \left[\theta^* \frac{R(1-\ell)(1-y)}{1-\lambda} + \frac{[\ell-\lambda]^+}{1-\lambda} - (\bar{R} + \bar{V}) \right] d\lambda \\ &\quad - \int_{\bar{\lambda}}^1 \frac{\ell + (1-\ell)\xi}{\lambda} (\bar{R} + \bar{V}) d\lambda = 0. \end{aligned} \tag{A.10}$$

$\bar{\Delta}^*$ is continuous in θ^* because all integrands are continuous and the discontinuity in v occurs only at one discrete point, $\hat{\lambda}$. From the existence of the upper and lower dominance regions and the continuity of $\bar{\Delta}^*$, there exists a θ^* such that $\bar{\Delta}^* = 0$ using the intermediate value theorem. Moreover, $d\bar{\Delta}^*/d\theta^* > 0$, so θ^* , which implies that the probability that a run occurs is unique and given by

$$\theta^* = \frac{\int_0^{\bar{\lambda}} (\bar{R} + \bar{V}) d\lambda + \int_{\bar{\lambda}}^1 \frac{\ell + (1-\ell)\xi}{\lambda} (\bar{R} + \bar{V}) d\lambda - \int_0^{\bar{\lambda}} \frac{[\ell-\lambda]^+}{1-\lambda} d\lambda}{\int_0^{\bar{\lambda}} \frac{R(1-\ell)(1-y)}{1-\lambda} d\lambda + \int_0^{\hat{\lambda}} V d\lambda}. \tag{A.11}$$

We now examine how a higher CBDC remuneration (or in the CBDC convenience yield), i.e., an increase in \bar{R} (or \bar{V}), affects run risk for different levels of liquidity transformation and payment services. From (A.11) we get that

$$\frac{d\theta^*}{d\bar{R}} = \frac{\int_0^{\bar{\lambda}} d\lambda + \int_{\bar{\lambda}}^1 \frac{\ell + (1-\ell)\xi}{\lambda} d\lambda}{\int_0^{\bar{\lambda}} \frac{R(1-\ell)(1-y)}{1-\lambda} d\lambda + \int_0^{\hat{\lambda}} V d\lambda} > 0, \tag{A.12}$$

which is also equal to $d\theta^*/d\bar{V}$. Equation (A.12) implies that increasing CBDC remuneration (or CBDC convenience) increases run risk irrespective of the degree of liquidity transformation and payment services of deposit. Intuitively, if the central bank supplies a superior means to store funds (or transact), then the risk of depositors withdrawing is higher, independent of the amount of liquidity that the

⁶³See Goldstein and Pauzner (2005), Infante and Vardoulakis (2021), and Kashyap, Tsomocos, and Vardoulakis (2024) for detailed mathematical steps getting from (A.9) to (A.10).

bank holds. However, the rate at which the run probability would increase with CBDC remuneration depends on the amount of liquid funds the bank holds. To understand this more precisely, we compute the following cross derivative:

$$\begin{aligned} \frac{d^2\theta^*}{d\bar{R}d\ell} &= \frac{\int_{\bar{\lambda}}^1 \frac{1-\xi}{\lambda} d\lambda}{\int_0^{\bar{\lambda}} \frac{R(1-\ell)(1-y)}{1-\lambda} d\lambda + \int_0^{\hat{\lambda}} V d\lambda} \\ &\quad - \frac{d\theta^*}{d\bar{R}} \frac{\int_{\bar{\lambda}}^1 \frac{R(1-\xi)/\xi}{1-\lambda} d\lambda + \frac{d\hat{\lambda}}{d\ell} V - \int_0^{\ell} \frac{R}{1-\lambda} d\lambda}{\int_0^{\bar{\lambda}} \frac{R(1-\ell)(1-y)}{1-\lambda} d\lambda + \int_0^{\hat{\lambda}} V d\lambda}, \end{aligned} \tag{A.13}$$

where $d\hat{\lambda}/d\ell = R(1-\xi)/(R-\xi) > 0$ from (A.3) and $d\theta^*/d\bar{R} > 0$ from (A.12) (similarly for $d^2\theta^*/(d\bar{V}d\ell)$). The cross derivative in (A.13) cannot be signed unambiguously due to opposing effects. On the one hand, higher ℓ reduces the need to inefficiently liquidate the risky asset as long as the number of withdrawals is less than what the bank can serve. On the other hand, a more liquid asset mix increases the probability of getting repaid conditional on the bank exhausting liquidity to serve all withdrawals, and, hence, increases the incentive to run. This trade-off is typical in bank run models with one-sided strategic complementarities akin to Goldstein and Pauzner (2005), and we perform numerical analysis to gauge the sign of the cross derivative. Table A.1 reports the change in the probability of a run as \bar{R} increases for different levels of ℓ . The left panel corresponds to lower ξ and the right panel to higher ξ . For both calibrations, $d^2\theta^*/(d\bar{R}d\ell)$ is negative moving across the columns, implying that lower liquidity transformation (higher ℓ) is associated with a *relative smaller* increase in run risk as \bar{R} increases. This result is intuitive but, as said, we can not preclude the possibility that the opposite holds under some alternative parameterizations.

To analyze whether V amplifies or mitigates the increase in run risk induced by higher CBDC remuneration, we compute the following cross derivative:

$$\frac{d^2\theta^*}{d\bar{R}dV} = - \frac{\int_0^{\hat{\lambda}} d\lambda}{\int_0^{\bar{\lambda}} \frac{R(1-\ell)(1-y)}{1-\lambda} d\lambda + \int_0^{\hat{\lambda}} V d\lambda} \frac{d\theta^*}{d\bar{R}} < 0. \tag{A.14}$$

Table A.1. Change in Run Risk for ℓ as \bar{R} Increases, Relative to $\bar{R} = 1$ for Different Levels of $\ell = 0.05$ and ξ

\bar{R}	$\xi = 0.7$						$\xi = 0.9$			
	$\ell = 0.05$	$\ell = 0.10$	$\ell = 0.15$	$\ell = 0.20$	$\ell = 0.20$	$\ell = 0.05$	$\ell = 0.10$	$\ell = 0.15$	$\ell = 0.20$	
1.00	—	—	—	—	—	—	—	—	—	
1.01	0.62%	0.60%	0.58%	0.57%	0.57%	0.44%	0.43%	0.43%	0.43%	
1.02	1.25%	1.20%	1.17%	1.14%	1.14%	0.88%	0.87%	0.86%	0.86%	
1.03	1.87%	1.80%	1.75%	1.71%	1.71%	1.31%	1.30%	1.29%	1.28%	

Note: $R = 3$, $V = 80$ bps, $\bar{V} = 0$.

This result is intuitive. Depositors forfeit the convenience yield when they withdraw their deposits for CBDC, which essentially reduces run risk all else being equal.⁶⁴

Finally, we consider how CBDC remuneration affects run risk if institutions are also allowed to invest ℓ in CBDC earning \bar{R} (but not \bar{V}). Adjusting (A.11) we get that

$$\frac{d\tilde{\theta}^*}{d\bar{R}} = \frac{\int_0^{\bar{\lambda}} d\lambda + \int_{\bar{\lambda}}^1 \frac{\ell\bar{R} + (1-\ell)\xi}{\lambda} d\lambda + \int_{\bar{\lambda}}^1 \frac{\ell}{\lambda} \bar{R} d\lambda - \int_0^{\ell\bar{R}} \frac{\ell}{1-\lambda} d\lambda}{\int_0^{\bar{\lambda}} \frac{R(1-\ell)(1-y)}{1-\lambda} d\lambda + \int_0^{\hat{\lambda}} V d\lambda} - \tilde{\theta}^* \frac{\int_{\ell\bar{R}}^{\bar{\lambda}} \frac{R\ell}{\xi(1-\lambda)} d\lambda + \frac{d\hat{\lambda}}{d\bar{R}} V}{\int_0^{\bar{\lambda}} \frac{R(1-\ell)(1-y)}{1-\lambda} d\lambda + \int_0^{\hat{\lambda}} V d\lambda}, \tag{A.15}$$

where $\tilde{\theta}^*$ is the run probability when institutions can also earn the CBDC remuneration. Adjusting (A.3) we also get $d\hat{\lambda}/d\bar{R} = \ell R / (R - \xi) > 0$. Compared to (A.12), (A.15) has two extra negative terms, mitigating the effect of higher \bar{R} on run risk, and one extra positive term, accruing from the one-sided complementarity property discussed earlier. Table A.2 shows the run risk for various ℓ and remuneration rates on CBDC. Comparing Table A.2 and Table A.1 shows that for sufficiently high ℓ , run risk is lower when institutions are allowed to hold CBDC and earn its remuneration rate (see, for example, $\ell = 0.2$, $\bar{R} = 1.03$, and $\xi = 0.9$). Hence, higher remuneration rates on CBDC mitigate run risk. For lower ℓ , the adverse effect of more liquidity in run incentives, described earlier, dominates and run risk may be higher (see, for example, $\ell = 0.05$, $\bar{R} = 1.03$, and $\xi = 0.7$).

⁶⁴Our assumption that payment services are enjoyed by depositors only if they are repaid at $t = 2$ in their entirety is important for this result.

Table A.2. Change in Run Risk for ℓ as \bar{R} Increases When Institutions Can Hold CBDC and Earn Its Remuneration Rate, Relative to $\bar{R} = 1$ for Different Levels of $\ell = 0.05$ and ξ

\bar{R}	$\xi = 0.7$						$\xi = 0.9$			
	$\ell = 0.05$	$\ell = 0.10$	$\ell = 0.15$	$\ell = 0.20$	$\ell = 0.05$	$\ell = 0.10$	$\ell = 0.15$	$\ell = 0.20$		
1.00	—	—	—	—	—	—	—	—		
1.01	0.63%	0.62%	0.60%	0.58%	0.44%	0.43%	0.43%	0.43%		
1.02	1.27%	1.23%	1.20%	1.16%	0.88%	0.87%	0.86%	0.84%		
1.03	1.90%	1.85%	1.80%	1.74%	1.32%	1.30%	1.29%	1.27%		

Note: $R = 3$, $V = 80$ bps, $\bar{V} = 0$.

References

- Ahnert, T., K. Assenmacher, P. Hoffmann, A. Leonello, C. Monnet, and D. Porcellacchia. 2024. "The Economics of Central Bank Digital Currency." *International Journal of Central Banking* 20 (4): 221–74.
- Ahnert, T., P. Hoffmann, A. Leonello, and D. Porcellacchia. 2023. "Central Bank Digital Currency and Financial Stability." Working Paper.
- Andolfatto, D. 2020. "Hot Money for a Cold Economy." Working Paper No. 2020-019, Federal Reserve Bank of St. Louis.
- Artavanis, N., D. Paravisini, C. Robles-Garcia, A. Seru, and M. Tsoutsoura. 2019. "Deposit Withdrawals." Working Paper.
- Azar, P., G. Baughman, F. Carapella, J. Gerszten, A. Lubis, J. P. Perez-Sangimino, D. Rappoport, C. Scotti, N. Swem, A. Vardoulakis, and A. Werman. 2022. "The Financial Stability Implications of Digital Assets." Finance and Economics Discussion Series No. 2022-058, Board of Governors of the Federal Reserve System.
- Bassett, W., M. Chosak, J. Driscoll, and E. Zakrajšek. 2014. "Changes in Bank Lending Standards and the Macroeconomy." *Journal of Monetary Economics* 62 (March): 23–40.
- Bassett, W., and D. Rappoport. 2022. "Enhancing Stress Tests by Adding Macroprudential Elements." Finance and Economics Discussion Series No. 2022-022, Board of Governors of the Federal Reserve System.
- Baughman, G., and F. Carapella. 2020. "Voluntary Reserve Targets." *Journal of Money, Credit and Banking* 52 (2–3): 583–612.
- . 2023. "A Simple Model of Voluntary Reserve Targets with Tolerance Bands." *Journal of Money, Credit and Banking* 55 (2–3): 655–72.
- Bindseil, U. 2020. "Tiered CBDC and the Financial System." ECB Working Paper No. 2351.
- Bindseil, U., F. Panetta, and I. Terol. 2021. "Central Bank Digital Currency: Functional Scope, Pricing and Controls." ECB Occasional Paper No. 2021-286.
- Brunnermeier, M. K., and D. Niepelt. 2019. "On the Equivalence of Private and Public Money." *Journal of Monetary Economics* 106 (October): 27–41.

- Calomiris, C. W., and C. M. Kahn. 1991. "The Role of Demandable Debt in Structuring Optimal Banking Arrangements." *American Economic Review* 81 (3): 497–513.
- Carapella, F., and J. Flemming. 2020. "Central Bank Digital Currency: A Literature Review." FEDS Notes, November 9, Board of Governors of the Federal Reserve System.
- Carlson, M., B. Duygan-Bump, F. Natalucci, B. Nelson, M. Ochoa, J. Stein, and S. Van den Heuvel. 2016. "The Demand for Short-Term, Safe Assets and Financial Stability: Some Evidence and Implications for Central Bank Policies." *International Journal of Central Banking* 12 (4): 307–33.
- Chiu, J., S. Davoodalhosseini, J. H. Jiang, and Y. Zhu. 2023. "Bank Market Power and Central Bank Digital Currency: Theory and Quantitative Assessment." *Journal of Political Economy* 131 (5): 1213–48.
- Chiu, J., and C. Monnet. 2024. "Public and Private Money Creation for Distributed Ledgers: Stablecoins, Tokenized Deposits, or Central Bank Digital Currencies?" Staff Working Paper No. 24-35, Bank of Canada.
- Cipriani, M., and G. La Spada. 2020. "Sophisticated and Unsophisticated Runs." Staff Report No. 956, Federal Reserve Bank of New York.
- Committee on Payments and Market Infrastructures. 2018. "Central Bank Digital Currencies." Bank for International Settlements.
- DeFusco, A., and A. Paciork. 2017. "The Interest Rate Elasticity of Mortgage Demand: Evidence from Bunching at the Conforming Loan Limit." *American Economic Journal: Economic Policy* 9 (1): 210–40.
- Duffie, D., and J. Younger. 2019. "Cyber Runs." Working Paper No. 51, Brookings Hutchins Center.
- Ennis, H., and T. Keister. 2009. "Bank Runs and Institutions: The Perils of Intervention." *American Economic Review* 99 (4): 1588–1607.
- Frost, J., L. Logan, A. Martin, P. McCabe, F. Natalucci, and J. Remache. 2015. "Overnight RRP Operations as a Monetary Policy Tool: Some Design Considerations." Finance and Economics Discussion Series No. 2015-010, Board of Governors of the Federal Reserve System.

- Goldstein, I., and A. Pauzner. 2005. "Demand-Deposit Contracts and the Probability of Bank Runs." *Journal of Finance* 60 (3): 1293–1327.
- Gorton, G., E. Klee, C. Ross, S. Ross, and A. Vardoulakis. 2026. "Leverage and Stablecoin Pegs." *Journal of Financial and Quantitative Analysis* 61 (1): 99–136.
- Gorton, G., and J. Zhang. 2023. "Taming Wildcat Stablecoins." *University of Chicago Law Review* 90 (3): 909–71.
- Greenwood, R., S. Hanson, and J. Stein. 2016. "The Federal Reserve's Balance Sheet as a Financial-Stability Tool." In *Designing Resilient Monetary Policy Frameworks for the Future*. Proceedings of the 2016 Economic Policy Symposium. Kansas City, MO: Federal Reserve Bank of Kansas City.
- Gross, D., and N. Souleles. 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics* 117 (1): 149–85.
- Huang, X., and T. Keister. 2024. "Can Redemption Fees Prevent Runs on Funds?" Available at <https://ssrn.com/abstract=4873088>.
- Infante, S., K. Kim, A. Orlik, and A. F. Silva. 2024. "Retail CBDC: Implications for Banking and Financial Stability." *Annual Review of Financial Economics* 16 (1): 207–32.
- Infante, S., K. Kim, A. Orlik, A. F. Silva, and R. J. Tetlow. 2022. "The Macroeconomic Implications of CBDC: A Review of the Literature." Finance and Economics Discussion Series No. 2022-076, Board of Governors of the Federal Reserve System.
- Infante, S., and A. P. Vardoulakis. 2021. "Collateral Runs." *Review of Financial Studies* 34 (6): 2949–92.
- Kashyap, A. K., R. Rajan, and J. C. Stein. 2002. "Banks as Liquidity Providers: An Explanation for the Coexistence of Lending and Deposit-Taking." *Journal of Finance* 57 (1): 33–73.
- Kashyap, A. K., D. P. Tsomocos, and A. P. Vardoulakis. 2024. "Optimal Bank Regulation in the Presence of Credit and Run Risk." *Journal of Political Economy* 132 (3): 772–823.
- Keister, T., and C. Monnet. 2022. "Central Bank Digital Currency: Stability and Information." *Journal of Economic Dynamics and Control* 142 (September): Article 104501.

- Keister, T., and D. Sanches. 2022. "Should Central Banks Issue Digital Currency?" *Review of Economic Studies* 90 (1): 404–31.
- Kim, M., D. Kilger, and B. Vale. 2003. "Estimating Switching Costs: The Case of Banking." *Journal of Financial Intermediation* 12 (1): 25–56.
- Kydland, F., and E. Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85 (3): 473–91.
- Lee, A. 2021. "What is Programmable Money?" FEDS Notes, June 23, Board of Governors of the Federal Reserve System.
- Li, L., Y. Li, M. Macchiavelli, and X. Zhou. 2021. "Liquidity Restrictions, Runs, and Central Bank Interventions: Evidence from Money Market Funds." *Review of Financial Studies* 34 (11): 5402–37.
- Li, L., P. Strahan, and S. Zhang. 2020. "Banks as Lenders of First Resort: Evidence from the COVID-19 Crisis." *Review of Corporate Finance Studies* 9 (3): 472–500.
- Logan, L. 2021. "Impact of Abundant Reserves on Money Markets and Policy Implementation." Speech delivered at SIFMA Webinar, via videoconference, April 15.
- Piazzesi, M., and M. Schneider. 2022. "Credit Lines, Bank Deposits or CBDC? Competition and Efficiency in Modern Payment Systems." Working Paper.
- Rose, J. 2023. "Understanding the Speed and Size of Bank Runs in Historical Comparison." *Economic Synopses* (Federal Reserve Bank of St. Louis) 12: 1–5.
- Rubinfeld, S. 2019. "Bangladesh Bank Sues Filipino Lender in U.S. Court Over Hack Heist: The New York Fed is Assisting the Bangladeshi Central Bank with the Lawsuit." *Wall Street Journal*, February 4.
- Schilling, L., J. Fernández-Villaverde, and H. Uhlig. 2024. "Central Bank Digital Currency: When Price and Bank Stability Collide." *Journal of Monetary Economics* 145 (July): Article 103554.
- Shy, O. 2002. "A Quick-and-Easy Method for Estimating Switching Costs." *International Journal of Industrial Organization* 20 (1): 71–87.
- Stenbacka, R., and T. Takalo. 2019. "Switching Costs and Financial Stability." *Journal of Financial Stability* 41 (April): 14–24.

- Van den Heuvel, S. 2019. “The Welfare Effects of Bank Liquidity and Capital Requirements.” 2019 Meeting Paper No. 325, Society for Economic Dynamics.
- Vissing-Jorgensen, A. 2021. “The Treasury Market in Spring 2020 and the Response of the Federal Reserve.” *Journal of Monetary Economics* 124 (November): 19–47.
- Whited, T., Y. Wu, and K. Xiao. 2022. “Central Bank Digital Currency and Banks.” Working Paper.