



INTERNATIONAL JOURNAL OF CENTRAL BANKING
Volume 22, Number 4
October 2026



INTERNATIONAL JOURNAL OF CENTRAL BANKING

(When) Do Banks React to Anticipated Capital Reliefs?

Guillaume Arnould, Benjamin Guin, Steven Ongena, and Paolo Siciliani

Understanding the Joint Dynamics of Inflation and Wage Growth in the Euro Area

Vahagn Galstyan

What Drives Euro-Area Financial Markets? The Role of U.S. Spillovers and Global Risk

Lennart Brandt, Arthur Saint Guilhem, Maximilian Schröder, and Ine Van Robays

Monetary Policy Strategies to Navigate Post-Pandemic Inflation: An Assessment Using the ECB's New Area-Wide Model

Matthieu Darracq Pariès, Antoine Kornprobst, and Romanos Priftis

ECB Euro Liquidity Lines

Silvia Albrizio, Iván Kataryniuk, Luis Molina, and Jan Schäfer

A Business Conditions Real-Time Measure of Output Gap Deviations

Alex Ilek and Yuval Mazar

"El Niño" and "La Niña": Revisiting the Impact on Food Commodity Prices and Euro-Area Consumer Prices

Fructuoso Borrallo, Lucía Cuadro-Sáez, Corinna Ghirelli, and Javier J. Pérez

Low for Long: Could Extended Periods of Ultra-low Monetary Policy Rates Have Harmful Effects?

Etibar Jafarov and Enrico Minnella



(When) Do Banks React to Anticipated Capital Reliefs? <i>Guillaume Arnould, Benjamin Guin, Steven Ongena, and Paolo Siciliani</i>	1
Understanding the Joint Dynamics of Inflation and Wage Growth in the Euro Area <i>Vahagn Galstyan</i>	45
What Drives Euro-Area Financial Markets? The Role of U.S. Spillovers and Global Risk <i>Lennart Brandt, Arthur Saint Guilhem, Maximilian Schröder, and Ine Van Robays</i>	97
Monetary Policy Strategies to Navigate Post-Pandemic Inflation: An Assessment Using the ECB's New Area-Wide Model <i>Matthieu Darracq Pariès, Antoine Kornprobst, and Romanos Priftis</i>	159
ECB Euro Liquidity Lines <i>Silvia Albrizio, Iván Kataryniuk, Luis Molina, and Jan Schäfer</i>	209
A Business Conditions Real-Time Measure of Output Gap Deviations <i>Alex Ilek and Yuval Mazar</i>	281
“El Niño” and “La Niña”: Revisiting the Impact on Food Commodity Prices and Euro-Area Consumer Prices <i>Fructuoso Borrallo, Lucía Cuadro-Sáez, Corinna Ghirelli, and Javier J. Pérez</i>	319
Low for Long: Could Extended Periods of Ultra-low Monetary Policy Rates Have Harmful Effects? <i>Etibar Jafarov and Enrico Minnella</i>	399

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**.

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, <i>Saudi Central Bank</i>	Jae Won Lee, <i>Bank of Korea</i>
Sergio Nicoletti Altimari, <i>Banca d'Italia</i>	Carlos Lenz, <i>Swiss National Bank</i>
Roc Armenter, <i>Federal Reserve Bank of Philadelphia</i>	Ye Liu, <i>People's Bank of China</i>
Kartik Athreya, <i>Federal Reserve Bank of New York</i>	Jochen Mankart, <i>Central Bank of Slovakia</i>
Marco Bardoscia, <i>Bank of England</i>	Karel Mertens, <i>Federal Reserve Bank of Dallas</i>
Jan Marc Berk, <i>The Nederlandsche Bank</i>	Andre Minella, <i>Central Bank of Brazil</i>
Paul Castillo, <i>Central Reserve Bank of Peru</i>	Gerard O'Reilly, <i>Central Bank of Ireland</i>
Laurent Clerc, <i>Bank of France</i>	Michael D. Patra, <i>Reserve Bank of India</i>
Paul Conway, <i>Reserve Bank of New Zealand</i>	Anna Paulson, <i>Federal Reserve Bank of Chicago</i>
Jose Gabriel Cuadra Garcia, <i>Bank of Mexico</i>	Thórarinn G. Pétursson, <i>Central Bank of Iceland</i>
Francisco C. Dakila Jr., <i>Bangko Sentral ng Pilipinas</i>	Andrea Raffo, <i>Federal Reserve Bank of Minneapolis</i>
Bernardus F. N. Doornik, <i>Central Bank of Brazil</i>	Trevor Reeve, <i>Federal Reserve Board</i>
Prof. Falko Fecht, <i>Deutsche Bundesbank</i>	Sigal Ribon, <i>Bank of Israel</i>
Marcus Fum, <i>Monetary Authority of Singapore</i>	Kasper Roszbach, <i>Norges Bank</i>
Carlos Garriga, <i>Federal Reserve Bank of St. Louis</i>	Krislert Samphantharak, <i>Bank of Thailand</i>
Joseph Gruber, <i>Federal Reserve Bank of Kansas City</i>	Çağrı Sarıkaya, <i>Central Bank of Turkey</i>
Philipp Hartmann, <i>European Central Bank</i>	Frank Smets, <i>Bank for International Settlements</i>
Sarah Hunter, <i>Reserve Bank of Australia</i>	Ulf Söderström, <i>Sveriges Riksbank</i>
Mugur Isărescu, <i>National Bank of Romania</i>	Salah-Eddine Taleb, <i>Bank of Algeria</i>
Esa Jokivuolle, <i>Bank of Finland</i>	George Tavlas, <i>Bank of Greece</i>
Edward Knotek, <i>Federal Reserve Bank of Cleveland</i>	Paula Tkac, <i>Federal Reserve Bank of Atlanta</i>
Anna Kovner, <i>Federal Reserve Bank of Richmond</i>	Dobieslaw Tymoczko, <i>National Bank of Poland</i>
Sharon Kozicki, <i>Bank of Canada</i>	Shingo Watanabe, <i>Bank of Japan</i>
Signe Krogstrup, <i>Danmarks Nationalbank</i>	Rafael Wouters, <i>National Bank of Belgium</i>
Ana Cristina Leal, <i>Bank of Portugal</i>	Egon Zakrajšek, <i>Federal Reserve Bank of Boston</i>
Sylvain Leduc, <i>Federal Reserve Bank of San Francisco</i>	

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

(When) Do Banks React to Anticipated Capital Reliefs?*

Guillaume Arnould,^a Benjamin Guin,^b Steven Ongena,^c
and Paolo Siciliani^b

^aInternational Monetary Fund

^bBank of England

^cUniversity of Zurich, Swiss Finance Institute, KU Leuven, and CEPR

We study how banks react to policy announcements during the representative policy cycle of consultation and policy publication. We use unique data covering the population of all mortgage transactions in the U.K., complemented with regulatory risk assessments of banks and their supervisory interactions. We show that banks likely to benefit from lower capital requirements increase the size of this capital relief by permanently investing into low-risk assets after the publication of the policy. Within this group, those banks with a higher frequency of supervisory interactions already react to the first step of policy development, the publication consultation. We show how these results can be used to estimate a lower bound on the cost of capital for smaller banks, for which such estimates are typically difficult to obtain.

JEL Code: G21.

*For comments and discussions, we would like to thank two anonymous referees, Thorsten Beck, Huberto Ennis (the editor), Chris Frank, Kathryn Inglis, Daniel Paravisini, and Neeltje van Horen. We also thank Jack Leslie for his early involvement in this project and Laavaniya Srithar for excellent research assistance. Any views expressed are solely those of the authors. They should not be taken to represent those of the Bank of England or as a statement of Bank of England policy. This paper should not be reported as representing the views of members of the Monetary Policy Committee, Prudential Regulation Committee or Financial Policy Committee. The views expressed in this paper are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management. Guillaume Arnould's contributions were completed whilst employed by the Bank of England. Steven Ongena acknowledges financial support from *ERC ADG 2016 - GA 740272 lending*.

1. Introduction

While the role of forward guidance in communicating monetary policy is well documented (McKay, Nakamura, and Steinsson 2016; Andrade et al. 2019), a significant gap remains in understanding how banks respond to anticipated regulatory changes communicated at various stages of the policy cycle. Identifying the effect of such announcements is crucial, as changes in minimum capital requirements affect bank lending (Gropp et al. 2019) and the real economy (Fraisse, Lé, and Thesmar 2020). This paper investigates a salient policy change that incentivizes banks to change their asset composition to reduce their capital requirements ahead of its firm-specific implementation, which occurs during their supervisory reviews. We examine whether banks anticipate and respond to potential capital requirement reliefs and identify which policy life-cycle announcements trigger these reactions.

The policy we examine is aimed at reducing minimum capital requirements for certain, eligible small banks that are deemed well-managed while leaving them unchanged for the remaining ones. The size of the reduction depends on the riskiness of their loan portfolios at the time of their supervisory review. In our setup, we follow announcements of the policy during the typical policy cycle: the publication of a consultation paper soliciting industry feedback on the proposed policy, and the publication of the final policy in the form of a policy statement. We also shed light on the role of supervisory oversight by exploring interactions between banks and their appointed supervisors.

We undertake a difference-in-differences analysis with multiple treatment dates in the spirit of Rodnyansky and Darmouni (2017) to study whether banks that were reasonably likely to be eligible reacted to these policy communications. Specifically, we examine whether they invested in low-risk mortgages to increase their expected capital relief, which we refer to as “prepositioning,” relative to other banks that were reasonably certain about their ineligibility. To this end, we use unique data on the population of all mortgage transactions in the United Kingdom between 2016 and 2019. The data come from the Financial Conduct Authority’s (FCA’s) Product Sales Database, which is well-established in the existing literature (Cloyne et al. 2019). We complement these data with bank balance

sheet information to also study banks' broader balance sheet adjustments beyond mortgages. Moreover, we use confidential information on banks' regulatory risk scores and the exact dates of their supervisory reviews combined with unique regulatory data on the frequency of bilateral interactions (i.e., meetings, site visits, and calls) between banks and their supervisors.

Our results suggest that banks that are likely to receive capital reliefs invest in low-risk mortgages prior to the supervisory review, which implements the policy at the bank. In turn, this prepositioning by banks leads to a reduction in the riskiness of their overall mortgage portfolio, which is persistent even after the implementation of the policy. By contrast, we find that banks that are unlikely to be eligible for capital relief tend to move in the opposite direction, increasing their exposure to riskier mortgages. This suggests that the two groups of banks are close competitors.

By the same token, we find evidence of a first-mover advantage in the sense that banks that have more time to adjust their asset composition fail to do so compared with banks that needed to react immediately after the publication of the policy statement. We also find evidence that banks with a high frequency of supervisory interactions already adjust the composition of their mortgage portfolios in response to the issuance of the consultation paper, whereas banks with a comparatively lower frequency of supervisory interaction wait until the publication of the policy statement. These results suggest that a higher frequency of supervisory interactions allows banks to better understand whether they are likely to be within the scope of the proposed policy.

This paper relates to three strands of the literature. First, we extend previous research on the announcement effects of bank regulation. The existing literature examines financial market reactions to new regulation (Schäfer, Schnabel, and di Mauro 2016; Bruno, Onali, and Schaeck 2018), but there is limited evidence on how banks themselves respond. Irani et al (2021) study U.S. banks' reactions to unexpected Basel III capital shortfalls, finding that they reduce their loan exposure, with nonbanks stepping in.¹ While Irani et al.

¹Similarly, Berrospide and Edge (2024) find that U.S. banks subject to stress testing react to the imposition of higher capital buffers by cutting their new

(2021) focus on *unexpected increases* in capital requirements, we analyze banks' reactions to future *reductions* in capital requirements. This distinction is crucial, as banks' responses to these changes may not be symmetric. Moreover, we can examine whether banks *anticipate* the final policy during its uncertain pre-implementation phase. Indeed, we find that eligible banks increase exposure to low-risk assets to secure expected capital reliefs, whereas Irani et al. (2021) document reduced loan exposures *ex post*.

Second, our paper contributes to the scarce literature on how regulatory uncertainty, specifically about the level of capital requirements, affects bank lending. It is well established that banks react to increased uncertainty about their ability to meet minimum capital requirements by holding (costly) voluntary capital buffers as a precaution (e.g., Jokipii and Milne 2008; Stolz and Wedow 2011; Valencia 2016), and that this often results in reduced lending (e.g., Buch et al. 2015; Kara and Yook 2019; Soto 2021).² In this literature, uncertainty entails a downside risk in the form of a costly breach of minimum capital requirements. By contrast, in our case regulatory uncertainty entails an upside risk for banks in the form of a reduction in capital requirements, which comes with lower capital costs. We complement the literature by finding that banks with more frequent supervisory interaction are more responsive when adjusting their asset composition. This behavior is consistent with the goal of maximizing the expected benefit of the policy proposal. It suggests that these banks are more confident about their eventual eligibility once the policy is finalized.

Third, our results are informative for banks' cost of capital considerations (Kisin and Manela 2016). For small banks, which are not listed on stock markets, such estimates typically do not exist. We address this gap by proposing a conceptual framework that showcases how our estimates can be used to gauge a lower bound of the implied cost of regulatory capital for such banks. We find that the

corporate loan originations, although this effect did not appear to affect borrowing firms' overall debt, investment, and employment, suggesting that firms can substitute for the reduction in loans from banks that participate in the stress tests.

²Eckley, Francis, and Mathur (2021) find similar results when the source of uncertainty is, specifically, the future level of capital requirements.

cost of capital for small banks is relatively insensitive to the size of the bank, and is comparable to large banks, at about 11 percent.

The remainder of the paper proceeds as follows. Section 2 sketches the background of the policy. Section 3 describes our data, while Section 4 discusses our empirical methodology. Section 5 presents the main results. Section 6 lays out an application with respect to the estimation of the shadow cost of capital for small banks. Section 7 concludes.

2. Background on the Policy

Under the Basel Committee's pillar 2 framework (BCBS 2001), supervisors can impose bank-specific capital requirements to capture risks that are not adequately, or not at all, captured under the standard pillar 1 minimum capital requirements, under which banks are required to hold capital resources of at least 8 percent of their risk-weighted asset bases. Credit risk typically represents the main component of banks' risk-weighted asset bases. Banks have two options to calculate risk weights: (i) they can estimate their own risk weights under the internal ratings-based (IRB) modeling approach, which is subject to explicit supervisory approval to ensure the integrity of these internal risk assessments; or (ii) they can adopt the standardized approach (SA), where risk weights are set in a standardized manner, supported by external credit assessments. The latter approach is typically adopted by smaller, less established, and less sophisticated banks.

In February 2017, the U.K. Prudential Regulation Authority (PRA), henceforth "the regulator," published a consultation paper proposing refinements to its approach to setting bank-specific capital requirements under the pillar 2 framework (PRA 2017). Specifically, the regulator sought to address concerns over higher risk weights under SA compared to those under IRB models. This was especially the case for residential mortgages with low loan-to-value (LTV) ratios, where SA risk weights can be up to seven times higher than under IRB, as shown in Table 1.

Under the proposals, the regulator would allow supervisors to exercise judgment and reduce bank-specific variable pillar 2A capital requirements for banks and building societies (henceforth, simply "banks") under the SA, where appropriate. In assessing capital

Table 1. IRB versus SA Gap in Risk Weights for Residential Mortgages

LTV Bands	Risk Weights under SA	Risk Weights under IRB
$0\% \leq \text{LTV} < 50\%$	35.0%	5.3%
$50\% \leq \text{LTV} < 60\%$	35.0%	9.1%
$60\% \leq \text{LTV} < 70\%$	35.0%	11.6%
$70\% \leq \text{LTV} < 80\%$	35.0%	16.6%
$80\% \leq \text{LTV} < 90\%$	36.0%	22.4%
$90\% \leq \text{LTV} < 100\%$	43.0%	33.3%

Note: Risk weights under IRB are average risk weights based on exposure amounts. They are calculated based on expected and unexpected losses prior to the implementation of the policy (PRA 2020). Subsequently, the gap between SA and IRB will have changed. Moreover, Basel 3.1 seeks to address that gap by lowering risk weights for low-LTV mortgages capitalized under the SA (Bank of England 2023).

needs, supervisors would consider the greater degree of conservatism that may apply to risk weights derived under the SA (Table 1, column 2) compared with those from IRB models for certain types of exposures, using so-called IRB credit risk benchmarks (Table 1, column 3). These are average risk weights for different types of exposure across banks that use IRB models. The aim of the policy was to address the concern that capital standards are overly prudent for banks using the SA for credit risk exposures. An excessive IRB versus SA gap might disadvantage SA banks in safer low-LTV bands, thus inducing them to increase their exposure in riskier high-LTV bands.

Repullo and Suarez (2004) argue that the IRB versus SA gap in risk weights could induce lenders to specialize in segments where they benefit from a comparative advantage (or, where the comparative disadvantage is reduced; see also Rime 2005; Ruthenberg and Landskroner 2008; Gropp, Hakenes, and Schnabel 2011). In this respect, Benetton et al. (2021) find that the gap in risk weights for U.K. residential mortgages between banks on IRB versus SA allows the former to undercut the latter, pushing SA banks to increase their exposure in riskier high-LTV segments. Building on this theoretical foundation and evidence base, the proposals were designed not only to mitigate competition distortions caused by the IRB versus SA

gap but also to address prudential risks posed by less sophisticated lenders increasing their exposure to higher-risk assets.

In October 2017, the regulator finalized this policy by publishing a policy statement, with an implementation date of January 1, 2018 (PRA 2017). Subsequently, this new approach was implemented as part of the Supervisory Review and Evaluation Process (SREP) reviews under the pillar 2 framework. Specifically, whereas the largest systemic banks are subject to this review process annually and simultaneously, smaller ones (i.e., those targeted by this policy) are reviewed on a rolling two- or three-year cycle (from 2018 to 2020), which is staggered over time. This proportionate approach also takes into consideration the limited supervisory resources available to oversee a large population of smaller banks.³

Under the SREP framework, banks are required to first undertake an internal capital adequacy assessment process (ICAAP) to assess on an ongoing basis their own capital adequacy depending on the level and nature of the risks to which they are or might be exposed. The regulator will then review the submitted ICAAP as part of the SREP. The ICAAP is based on the latest available annualized financial statement data, typically from the end of the year that precedes the SREP.

The calculation of the reduction in the pillar 2A capital framework does not follow mechanically from the comparison with IRB benchmarks. Supervisors maintain full discretion informed by a business model analysis, including consideration of whether the benchmark is representative of bank-specific risk profiles. Banks considered to be relatively low-risk and well-managed would be considered eligible for capital reliefs. By contrast, banks that do not ensure sound management and coverage of their risks or that have concentrated exposures to asset classes for which the regulator does not have sufficient data to produce a reliable IRB benchmark would be considered ineligible. In addition, the published eligibility criteria of the policy were left relatively flexible to allow scope for supervisory judgment. As a result, banks on the SA faced uncertainty as to whether they were eligible.

³It is worth noting that the PRA expects banks to publicly disclose the amount of total capital requirements (given by the sum of pillar 1 and pillar 2 requirements) which apply to them. See PRA (2019).

In conclusion, one aim of this policy was to incentivize SA banks to rebalance their portfolio from riskier high-LTV residential mortgages toward safer low-LTV ones. This was deemed to be beneficial not only in terms of competition, by leveling the playing field between IRB and SA banks (Benetton et al. 2021), but also with regard to the safety and soundness of these relatively less sophisticated SA banks.

Given that the calculation of the reduction in capital requirements was based on the portfolio allocation reported by banks as part of their ICAAP preceding the SREP, banks had the opportunity to increase their exposure to low-LTV residential mortgages before the actual policy implementation (a behavior we label as “prepositioning”) in order to increase not only the potential benefits from the implementation in terms of capital requirements, but also the likelihood of being considered eligible in the first place. This possible change in viewpoint is because the resulting safer portfolio allocation would in turn ameliorate the overall risk profile of the bank in question, which is one of the eligibility criteria outlined in the consultation paper and confirmed in the policy statement.

For example, the SREP of a bank in May 2019 would have used balance sheet information from its ICAAP based on data typically from the end of the previous year. Therefore, the bank would have had the opportunity to increase its exposure to low-LTV mortgages during 2018. In doing so, however, the bank would have had to depart from what it considered to be the best portfolio allocation given the current capital requirements. As lower-LTV mortgages are typically priced at lower rates than high-LTV mortgages, prepositioning in low-LTV mortgages would entail an opportunity cost, arising from adjusting the current allocation to that which would maximize the expected benefit from the policy implementation. This benefit is a function of the probability of being eligible and the magnitude of capital relief.

3. Data

3.1 Bank Eligibility Criteria

Both the consultation paper and the policy statement highlight the criteria that determine whether or not a bank is eligible for capital

relief under the policy. Specifically, the consultation paper mentions three criteria that are relevant for banks' eligibility considerations: (i) banks should have sound management, (ii) they should have a low risk profile, and (iii) they should have no concentrated exposure to asset classes for which the regulator does not have an IRB benchmark.⁴ In addition, banks that use the IRB approach to capitalize credit risk are not eligible. We translate these criteria into metrics that allow us to build a list of banks that are likely to consider themselves as eligible.

Criterion 1: Sound management. We use the list of banks with a risk management and governance (RM&G) scalar to incorporate these criteria.⁵ We assume that a bank will consider itself as having sound management if the regulator did not flag its risk management practice as bad. In other words, the baseline assumption of our approach is that all banks have sound management. Supervisors may flag significantly weak RM&G for banks whose capital requirements are increased as a result via an RM&G scalar.⁶ Hence, we argue that banks that have an RM&G scalar will consider themselves as not eligible for the policy.

Criterion 2: Low risk profile. Supervisors rank banks over a cardinal scale based on several risk drivers (i.e., from low- to high-risk).⁷ Those risk scores are for internal purposes only and, thus, are not communicated to banks. However, banks have regular communication with their supervisors, which culminates in a holistic assessment of their specific risk profile. This feedback can help banks understand how risky their supervisors perceive them to be and, thus, whether they are likely to be eligible. The idea is that a bank below a certain threshold should consider itself as being perceived as a low-risk bank by its supervisors. Specifically, using risk

⁴The policy applied to all asset categories that could give rise to a gap in risk weights between SA and IRB, although by far the most prominent of these was residential mortgages, which is the focus of our analysis.

⁵Under the European Banking Authority's SREP guidelines, competent authorities should set additional own-funds requirements to cover risks posed by RM&G weaknesses where appropriate as an interim measure, while the deficiencies are addressed; see EBA (2014).

⁶For more details on the methodology adopted, see PRA (2023, paragraphs 60–64).

⁷For more details on the internal risk scoring methodology, see PRA (2018a, Chapter 3) and also Bholat et al. (2017).

scores from June 2017, we build two proxies for the true criterion to construct the list of eligible banks: (i) no risk score above two-thirds of the maximum score, and (ii) an average risk score less than one-half of the maximum score.

Both proxies rely on the idea that the safest banks can gauge whether they are considered as relatively low-risk. However, the first proxy is more precise and, arguably, less demanding for banks to be able to infer their eligibility, in the sense that in the absence of supervisory feedback focused on specific issues they can infer themselves to be considered low-risk. Indeed, the number of eligible banks falls from nine under criterion (i) to three under criterion (ii), so that only the most conservative banks are considered eligible under the latter proxy.

Supervisors provide both formal and informal communication to banks, which helps them gauge their eventual eligibility. To determine if banks with at least one high risk score are likely aware of their risk and eligibility, we also reviewed a sample of formal letters which supervisors sent to banks at the end of their SREP. We examined if the tone of the letter regarding a specific risk area changed once supervisors had given a high score. We then compared the tone of these letters with letters from banks that are not considered to be high-risk in the same areas. Our analysis found that letters to banks with high risk scores clearly signal supervisors' concerns about the specific area in question.⁸ Accordingly, we adopt the first criterion based on no risk score above two-thirds of the maximum score as our main approach. Additionally, we retain the second approach, which is based on a low average risk score, as an alternative proxy for robustness analysis.

Criterion 3: Exposure to assets with no IRB benchmark.

We do not capture this criterion explicitly, as we restrict our analysis to banks that only have some material exposures to residential mortgages; that is, with no fewer than 500 mortgages originated

⁸Table A.1 in the Appendix illustrates these differences in wording. Letters to banks with low risk scores tend to have a more positive tone (suggesting ways to improve), while banks with high risk scores have a more direct tone in their letter, pointing out the area they must improve upon. See Bholat et al. (2017) for a more in-depth analysis of the semantic content of those supervisory letters.

over the sample period.⁹ The consultation paper includes a list of asset classes of credit risk IRB benchmarks. We focus our analysis on residential mortgages, which are differentiated by LTV. We do this for two reasons: first, residential mortgages constitute the most material asset class for most retail banks; second, this is the asset class for which the gap in risk weights between the SA and the IRB approach is the widest.

Criterion 4: No IRB modeling. We restrict our sample to banks using the SA. The target of the policy is to reduce the gap in risk weights between banks using the SA and banks using IRB models. Banks that have an IRB license could be eligible for the policy if they use the SA for a fraction of their portfolio. However, it is unlikely that banks with an IRB license use SA for their mortgage portfolio given that risk weights are notably higher under SA. Therefore, we consider that banks that have an IRB license (at the entity level or at the group level) always use the IRB approach for their mortgage portfolio and will not deem themselves eligible. In addition, we also exclude these IRB banks from the control group of ineligible banks. This is because, compared to the type of banks likely to be eligible, banks using IRB models tend to be very large and with a significant competitive advantage in terms of cost of funding (FCA 2018). Therefore, for the sake of comparability, we restrict both our treatment and control groups to banks under the SA.

Table 2 summarizes the eligibility criteria set out in the consultation paper and confirmed with the policy statement, and how we take them into account in our analyses. Out of 47 banks in our sample, we classify 30 as being eligible based on our first approach to proxy for the risk profile of banks, while only 8 satisfied the alternative approach based on the average risk score criterion.

We complement this information with the exact date of the supervisory review (SREP) for each of the banks we include in our sample. Banks differ in the year of their initial SREP and in the length of the cycle until their next SREP: while most banks in our

⁹ Accordingly, we excluded three banks with fewer than 500 residential mortgage originations over the sample period. The inclusion of these small lenders in the sample does not alter the results.

Table 2. Eligibility Criteria—Consultation Paper

Eligibility Criteria from CP	Our Proxy
Sound Management	Banks with an RM&G scalar are not eligible
Low Risk Profile	Banks with no risk score above two-thirds of the maximum score are eligible (alternative criterion: an average risk score less than one-half of the maximum score)
No Concentrated Exposure to Assets w/o Benchmark	We restrict our analysis to banks with substantial mortgage exposure as a robustness test
Using the Standardized Approach	Banks using the internal ratings-based approach are not eligible

sample are on a three-year cycle, there are a few that are on a two-year or one-year cycle. To create clearly defined cohorts of banks, we reattribute the latter to the following SREP cohorts: e.g., a bank on a two-year cycle with a SREP in 2018 is also added to the 2020 cohort.

In our empirical analyses, we focus on the sample of 34 banks which had their SREP in 2019 and 2020, and not on the previous cohorts.¹⁰ Regarding the cohort of banks with a SREP in 2018, the publication of the policy statement in the fourth quarter of 2017 was too late for them to properly preposition by the end of the same year (i.e., the cutoff date for the collection of annual data to be used in the ICAAP submission). In summary, there are 13 banks that had their SREP in 2019, of which 9 were deemed to be eligible (under the criterion that the bank has no risk score above two-thirds of the maximum score);¹¹ and 21 had their SREP in 2020, with 10 eligible. Nevertheless, as explained in Section 5.4, our main empirical analyses and results are focused on the cohort of banks with a SREP in 2019.

¹⁰Therefore, we drop 13 banks on a three-year cycle with a SREP in 2018. The remaining banks in the 2018 cohort and all the banks with a SREP in 2017 are reallocated to the 2019 and 2020 cohorts according to the length of their respective cycles.

¹¹Under the alternative criterion (average risk score), there are 13 banks that had their SREP in 2019, of which 3 were deemed to be eligible; and 21 with a SREP in 2020, 2 of which were eligible.

Table 3. Differences in Bank Characteristics (SREP in 2019)

Variable	Eligible	Not Eligible	Difference	ND
LTV (%)	55.74	64.99	−9.24	−0.96
Mortgages (% of Total Assets)	68.47	55.12	13.35	0.48
Deposits (% of Total Assets)	84.01	79.70	4.30	0.24
Capital (% of RWA)	18.27	21.85	−3.58	0.37
Total Assets (£m)	1,741.27	1,264.02	477.24	0.25

Note: This table shows mean differences in key bank characteristics at the beginning of 2017 for banks with SREPs in 2019, by eligibility status. The third column shows the mean differences, and the last column presents normalized differences (NDs). ND measures the standardized difference between groups, with a critical threshold of 0.25 suggesting substantial differences. The data come from an updated version of the Bank of England’s HBRD database as described in de Ramon, Francis, and Milonas (2017), except for the LTV data, which are from the FCA’s Product Sales Database.

In Table 3, we compare the characteristics of 2019 SREP banks deemed eligible relative to banks classified as ineligible in the spirit of Gropp et al. (2019). Specifically, the table shows key bank characteristics before the publication of the consultation paper, and it calculates differences and normalized differences between these two groups.

The analysis reveals that eligible and ineligible banks are similar in size and deposit reliance. Eligible banks had larger total assets (mean of GBP 1.7 billion) compared to ineligible banks (mean of GBP 1.3 billion), with a normalized difference (ND) of around 0.25, which is right at the critical threshold of 0.25 suggested by Imbens and Rubin (2015) and Imbens and Wooldridge (2009). They also had a slightly higher deposit ratio (84 percent versus 80 percent), with an ND of 0.24, indicating a notable, though moderate, difference in their funding structures. Differences between these two groups are more pronounced in terms of mortgage portfolios and capital adequacy. Eligible banks hold a higher percentage of mortgages than ineligible ones (68 percent versus 55 percent) and have a lower average LTV ratio of new mortgages (56 percent versus 65 percent), with NDs of 0.48 and −0.96, respectively. Additionally, eligible banks have a lower capital ratio than ineligible ones

Table 4. Summary Statistics

	Mean	St. Dev.	Min.	Max.	Observations
LTV					
LTV	66.43	24.65	6.33	95.00	62,977
LTV (<80%)	0.60	0.49	0.00	1.00	62,977
LTV (<60%)	0.34	0.47	0.00	1.00	62,977
LTV (<50%)	0.26	0.44	0.00	1.00	62,977
Control Variables					
Refinancing	0.43	0.50	0.00	1.00	62,977
Fixed Rate	0.86	0.35	0.00	1.00	62,977
Short Incentive Period	0.42	0.49	0.00	1.00	62,977
Note: This table shows the summary statistics of the variables displayed in the regression tables. The sample includes observations from 2016:Q1 through 2019:Q4. LTV is winsorized at 1 percent and 99 percent to account for outliers.					

(18 percent versus 22 percent) with an ND of 0.37, reflecting significant differences in both risk profiles and capital strength.

3.2 Mortgage Transactions

For each of the 13 banks with a SREP in 2019, we match mortgage transactions of all mortgages granted from the first quarter of 2016 until the fourth quarter of 2019. These data come from the FCA's Product Sales Database. In this data set, we observe new mortgages and mortgages refinanced with a new bank ("switchers"). It includes the exact quarter of each of these transactions as well as further mortgage and borrower characteristics. Table 4 shows these characteristics in our sample. We exclude the two quarters of the publication of the consultation paper and policy statement, as these quarters are not clearly attributable to the time before publication or after. Our final sample consists of 62,977 mortgage transactions.

Most importantly for our analyses, our sample includes the LTV ratio of each mortgage, which is 66 percent on average (Table 4, column 1, row 1).¹² We use this variable to create a dummy variable

¹²The distribution of LTV is winsorized at 1 percent and 99 percent to account for outliers.

denoting whether a mortgage is below a certain LTV threshold or not. In our main analyses, we classify mortgages with an LTV below 60 percent as being at low LTV. Below this threshold, IRB risk weights are less than one-third of the corresponding weights under SA (Table 1). Thus, they entail a material disadvantage for banks using SA. This threshold also identifies the lowest pricing band conventionally used by banks (i.e., to qualify for the lowest available rates). About one-third of all mortgage transactions display an LTV of less than 60 percent (Table 4, column 1, row 3). We also adopt two alternative thresholds: (i) LTV below 80 percent, which is more conservative; and (ii) LTV below 50 percent, where the gap in risk weights is at its widest (Table 1, row 1). About 60 percent of all mortgages show LTVs of below 80 percent (Table 4, column 1, row 2); about one-fourth have an LTV below 50 percent (Table 4, column 1, row 4).

The Product Sales Database also provides a set of other relevant mortgage contract terms. Regarding the purpose of the mortgage transaction, on average about 43 percent of all mortgages are taken by refinancers (Table 4, column 1). Most mortgages in our sample are priced at a fixed rate (86 percent—Table 4, column 1, penultimate row). Finally, with respect to the length of the initial incentive period,¹³ 42 percent of all mortgage transactions have incentive periods of two years or shorter (Table 4, column 1, last row).

3.3 Frequency of Bilateral Supervisory Interaction

We complement this data set with information on the number of interactions between banks and their supervisors. As the number of interactions can vary substantially across banks and years, we do not use the absolute number of bank-supervisor interactions. Instead, we employ a measure of deviation from the median number of interactions within the four supervisory categories of banks. For each of these categories, we then estimate the median of the number of interactions in 2018 or 2019.¹⁴ We use the median instead

¹³In the U.K., mortgages are on a teaser rate, which typically lasts two years. After this period, the mortgage rate reverts to a relatively higher standard variable rate (SVR).

¹⁴We use the average of the median only for 2018 and 2019 given that 2017 is incomplete.

of the mean as it is robust to outliers of banks with very frequent interactions.

We consider a bank to have a high number of supervisory interactions if it met with its supervisor more than 1.5 times the median for its bank category. We chose the threshold of 1.5 times the median to capture only banks that had significantly more interactions with their supervisors than the remaining banks of their category. We create a dummy variable that assigns the value of one for banks with high supervisory interactions (and zero otherwise). In our sample of 13 banks with a supervisory review in 2019, we observe that 5 banks have a high number of supervisory interactions.

4. Empirical Strategy

Our empirical strategy follows Rodnyansky and Darmouni (2017). They use a difference-in-differences setup with multiple treatment dates to examine the effects of quantitative easing on bank lending behavior. In our main specification, we estimate the following regression model using ordinary least squares (OLS):

$$\begin{aligned} \text{Low } LTV_{i,b,t} = & \beta_1 \times \text{Eligible}_b \times \text{Post } PS_t + \beta_2 \times \text{Eligible}_b \\ & \times \text{Post } CP_t + \beta_3 \times \text{Post } PS_t + \beta_4 \times \text{Post } CP_t + \gamma_b + X_{i,b,t} + \varepsilon_{i,b,t}, \end{aligned}$$

where *Low LTV*_{*i,b,t*} is an indicator variable which equals one if the mortgage *i* by bank *b* at point in quarter *t* has an LTV ratio of less than 60 percent and zero otherwise. In other specifications, we use alternative thresholds and contrast them with a continuous measure of LTV. The comparison with a continuous measure of LTV is especially relevant from a policy evaluation perspective, in that it would be an unintended policy consequence if eligible banks increased their exposure at very high LTV (i.e., where interest rates are higher) to compensate for the increase at low LTV (i.e., polarized portfolio allocation), which might leave the average LTV unchanged. Therefore, the result desired from a policy perspective would be an increase at low LTV alongside a reduction in the continuous measure of LTV.

*Post CP*_{*t*} is a time dummy variable which indicates whether the mortgage transaction was made in the quarters following the consultation paper publication, i.e., in the second quarter of 2017 or

later (the variable equals zero otherwise). $Post\ PS_t$ is an indicator of whether the mortgage transaction was made in the quarters after the policy statement publication, i.e., in the first quarter of 2018 or later (zero otherwise). The coefficient of this variable can be interpreted as the marginal effect of the publication of the policy statement net of the effect of the publication of the consultation paper. We exclude the quarters of the two treatment dates, consultation paper (2017:Q1) and policy statement (2017:Q4), as these quarters include days before as well as days after the publications. We interact the time dummy variables with $Eligible_b$, which indicates whether a 2019 SREP bank b was deemed eligible, using our “no risk score above two-thirds” criterion in our baseline analyses. γ_b is the vector of bank fixed effects which controls for time-invariant bank characteristics. This dummy variable absorbs the main effect of $Eligible_b$, which does not appear in our regression equation. $X_{i,b,t}$ is a vector of contract characteristics of mortgage transaction i (incensive period as well as the type of interest rate, i.e., fixed rate). It also includes dummy variables to account for the geographic variation in the 12 U.K. regions included in our sample. We also interact these regional fixed effects with time.¹⁵

In addition, in an alternative specification, we interact the time dummy variables and the eligibility dummy variables with a third dummy variable, denoted as $Supervision(high)_b$. It indicates whether bank b was subject to a high frequency of bilateral supervisory interaction. The aim of this triple-interacted term is to uncover whether eligible banks with a higher degree of supervisory interaction responded more promptly to the publications of the consultation paper and the policy statement (especially the former) than eligible banks with a lower degree of supervisory interaction.¹⁶

Given the level of administration of the assessed policy, we cluster standard errors at the region and time level. Our sample includes observations in the four years from 2016:Q1 through 2019:Q4. Hence,

¹⁵In unreported robustness checks, we also control for refinancing $\times$ time fixed effects. We conduct two tests, both replacing the region $\times$ time fixed effects and adding to the region $\times$ time fixed effects. The results stay qualitatively very similar.

¹⁶We limit the estimation of this alternative specification to the cohort of banks with SREPs in 2019, as those in the 2020 cohort did not face the same degree of urgency to preposition in response to the consultation paper.

our estimated coefficients for the consultation paper ($Post\ CP_t$) and its interactions with $Eligible_b$ and $Supervision(high)_b$ can be interpreted as the treatment effects relative to the pre-event period from 2016:Q1 through 2016:Q4. Similarly, our estimated coefficients for the policy statement ($Post\ PS_t$) and its interactions with $Eligible_b$ and $Supervision(high)_b$ can be interpreted as additional treatment effects relative to the effect of the consultation paper.

Like Rodnyansky and Darmouni (2017), we also estimate the effect of each of the time quarters included in our sample relative to 2016:Q1 through 2016:Q3. Undertaking these quarter-by-quarter analyses is relevant for two reasons. First, it allows us to examine the validity of the parallel trends assumption of our difference-in-differences design. Specifically, we can examine whether eligible banks had started prepositioning before the publication of the consultation paper. Second, it allows us to examine the persistence of the effects over time. This is important from a policy evaluation perspective, in that a reversal of the prepositioning effect after the date relevant to the submission of the ICAAP (i.e., at the end of 2018 for the cohort of banks with a SREP in 2019) would raise suspicions of “window-dressing.” By contrast, the persistence of the prepositioning effect would be in line with the aim of the policy. Regarding the second reason, there might be a preexisting trend in our data with eligible banks investing in low-LTV mortgages prior to the announcement of the policy. This would invalidate the parallel trends assumption in our difference-in-differences design. Formally, we estimate the following regression model using OLS:

$$\begin{aligned} Low\ LTV_{i,b,t} = & \sum_{t=-1}^T \beta_t \times Eligible_b \times Time_t + \sum_{t=-1}^T \gamma_t \times Time_t \\ & + \gamma_b + X_{i,b,t} + \varepsilon_{i,b,t}, \end{aligned}$$

where $Time_t$ is an indicator of whether the observation is in quarter t (zero otherwise). For example, $Time_{t-1}$ indicates whether the mortgage transaction was done in 2016:Q4, the quarter before the publication of the consultation paper (zero otherwise). In the absence of prepositioning before the publication of the consultation paper, we expect the estimate of β_{-1} not to be statistically significantly different from zero. If it was negative, there might be evidence of

**Table 5. Difference in Average LTV
across Eligibility and over Time**

	Before (1)	Post (CP) (2)	Post (PS) (3)	Difference (4)
Eligible	66.743	69.029	61.744	−4.999***
Noneligible	64.498	63.494	73.728	9.229***
Note: This table shows the results of mean LTV across bank eligibility and time. It shows the results for banks that had their SREP in 2019 as well as banks with SREPs in 2017 on a one-/two-year cycle and banks with SREPs in 2018 on a one-year cycle. The first column shows average LTV in the period before the CP was published. The second column shows average LTV in the period after the CP was published and before the PS was published. The third column shows average LTV in the period after the PS was published. The last column tests differences in mean LTV between the first and the third column. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 level, respectively.				

prepositioning. Compared to the previous regression equation, we now include the quarters of the two treatment dates, consultation paper (2017:Q1) and policy statement (2017:Q4). All coefficients are estimated relative to 2016:Q1, as it allows us to estimate the effect of the three last quarters before the publication, from 2016:Q2 through 2016:Q4.

5. Results

5.1 The Effect of Capital Reliefs on Prepositioning

In Table 5, we examine whether banks with a SREP review in 2019 prepositioned in low LTV after the two event dates, i.e., the publication of the consultation paper (CP) in February 2017 and the publication of the policy statement (PS) in October 2017. The table presents simple differences in mean LTV, where column 1 shows average LTV before the publication of the CP, column 2 shows average LTV after the publication of the CP and before the publication of the PS, and column 3 shows average LTV after the publication of the PS. Column 4 tests the mean differences between LTV averages before any publication and mean LTV after the publication of the PS. In summary, starting from broadly similar levels, eligible banks reduced their average LTV after the publication of the PS, whereas

noneligible banks moved in the opposite direction. Both results are statistically significant at the 1 percent level.

To further explore this apparent incongruency, Table 6 presents the results of a linear regression model as outlined above. In columns 1–2, the dependent variable is banks' propensity to grant a low-LTV mortgage with an LTV of less than 60 percent (conditional on granting a mortgage). In columns 3–4, we show the effects by changing the dependent variable to a continuous measure of LTV (again conditional on granting a mortgage). Explanatory variables are indicators of bank eligibility (using the “below two-thirds” criterion) as well as two time dummy variables, which correspond to the quarters following the two policy publications.

The prepositioning effect is pronounced in the aftermath of the publication of the PS. In fact, following the publication of the PS, eligible banks increased their propensity to invest in a low-LTV mortgage (<60 percent) by almost 14 percentage points relative to ineligible banks, and this result is statistically significant at the 1 percent significance level (column 1). This estimated effect decreases slightly (by 1 percentage point) once we control for region $\times$ time fixed effects, but it remains statistically significant at the 1 percent significance level (column 2).

Results in columns 3–4 suggest that not only did the share of low-LTV mortgages increase, but also the overall LTV composition of new mortgages decreased. That is, banks did not offset lending into low-LTV mortgages by simultaneously increasing the share of lending to very high-LTV borrowers. In fact, following the publication of the PS, eligible banks decreased the average LTV of new mortgages by 7 percentage points relative to ineligible banks. These estimated effects are statistically significant at the 1 percent significance level. This result is consistent with the aim of the policy, as eligible banks also materially reduced the LTV level across their residential mortgage books.

The sharpness of these results is strengthened because the coefficient of the noninteracted term, *Post (PS)*, which captures the reaction of ineligible banks, is statistically significant at the 5 percent significance level (columns 1 and 3), but with an opposite sign. This contrast indicates that ineligible banks reacted to the prepositioning of eligible banks following the publication of the PS by adjusting their portfolios in the opposite direction, thus suggesting

Table 6. The Average Effect of Capital Reliefs on Prepositioning

Dependent Variable	LTV (<60 percent)		LTV	
	(1)	(2)	(3)	(4)
Eligible × Post (CP)	−0.056 (0.028)	−0.031 (0.015)	2.168 (1.211)	1.643 (0.789)
Eligible × Post (PS)	0.138*** (0.009)	0.127*** (0.008)	−7.227*** (0.503)	−7.076*** (0.424)
Post (CP)	0.066 (0.036)		−2.305 (1.602)	
Post (PS)	−0.122** (0.015)		5.702** (0.636)	
Bank FE	YES	YES	YES	YES
Region FE	YES	NO	YES	NO
Region × Time FE	NO	YES	NO	YES
Contract Variables	YES	YES	YES	YES
Observations	62,977	62,977	62,977	62,977
Number of Banks	13	13	13	13
Eligible Banks	9	9	9	9
Mean of Dep. Variable	0.344	0.344	66.433	66.433
R2	0.269	0.271	0.380	0.382
Clustered SE	YES	YES	YES	YES
Method	OLS	OLS	OLS	OLS

Note: This table shows the results of a linear regression model estimated using OLS on the mortgage transaction level. In columns 1 and 2, the dependent variable is the propensity to grant a low-LTV loan (below 60 percent). In columns 3 and 4, the dependent variable is the continuous variable, LTV. In all columns, this table shows the results for banks that had their SREP in 2019 as well as banks with SREPs in 2017 on a one-/two-year cycle and banks with SREPs in 2018 on a one-year cycle. Explanatory variables are a time dummy variable as well as its interaction with bank eligibility. In all columns, we control for contract variables and bank fixed effects. In columns 1 and 3, we include region fixed effects. In columns 2 and 4, we include region × time fixed effects. Contract variables include fixed rate, incentive period, and refinancing. Standard errors are clustered on the region and time level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 level, respectively.

that the two sets of banks are close competitors in the sense that ineligible banks are forced to accommodate the expansion of eligible banks in low-LTV segments by repositioning toward higher-LTV segments. This “crowding-out” effect could be an unintended consequence of the policy in question, given that two of the eligibility criteria are that banks should have sound management and have a low risk profile. In this sense, it is interesting to note that these effects are strongly reminiscent of what was theorized by Repullo and Suarez (2004) and corroborated by Benetton et al. (2021) with respect to the IRB versus SA gap.

There is no statistically significant effect in response to the publication of the CP in February 2017. This is somewhat remarkable as it was very likely that the policy proposal consulted upon would be ultimately confirmed and thus implemented in January 2018. This is because the U.K. regulator had already signaled its commitment to reduce the IRB versus SA gap in risk weights for residential mortgages in 2016 by reviewing the methodology for setting pillar 2A capital requirements.¹⁷ Therefore, this suggests that regulated banks wait for the finalization of a policy before reacting to it, no matter how well-anticipated the policy in question may be. Another factor that could explain this delayed reaction from banks is the uncertainty regarding their eventual eligibility under the proposed scoping criteria.

¹⁷On October 28, 2016, Sam Woods (Deputy Governor, Prudential Regulation and Chief Executive Officer, Prudential Regulation Authority), said: “I have long been troubled, as I know have members of the Treasury Committee, by the gap in risk-weights for low-LTV mortgages between banks who use models to calculate them and (typically smaller) banks who use a standardised, and therefore relatively crude, weighting provided by regulators. Now of course the leverage ratio is an essential complement to the risk-weighted framework which mitigates the effect of such disparities. Nevertheless they still bother me both in light of our secondary competition objective, given the risk of an un-level playing field, but importantly also in light of our safety and soundness objective because of the economic incentive it provides for standardised banks to concentrate on higher-LTV lending. [...] I intend that we will bring forward proposals under our Pillar 2 regime which should also reduce the risk that our capital standards are overly prudent for smaller banks using the standardised approach to credit risk to calculate their requirements – essentially by looking at capital requirements in the round rather than assuming that a simple “sum of the parts” approach will necessarily deliver the right answer.” (Woods 2016, p. 8)

5.2 *The Persistence of Effects of Capital Reliefs on Prepositioning*

In this section, we analyze the persistence of prepositioning by eligible banks. By doing so, we examine the validity of the parallel trends assumption of our difference-in-differences design. To that end, we estimate the effect of each of the time quarters included in our sample relative to the period from 2016:Q1 through 2016:Q3, following the approach of Rodnyansky and Darmouni (2017).


In Figure 1, we display the estimated quarterly coefficients and 95 percent confidence intervals using LTV (<60 percent) as our dependent variable. The estimated coefficient in 2016:Q4 is close to zero and statistically insignificant. This indicates that there is no evidence that eligible banks prepositioned in the quarter before the publication of the CP. In line with our previous results, we do not see evidence of prepositioning in the quarters following the publication of the CP in 2017:Q1. However, banks start prepositioning in the quarter following the publication of the PS, 2018:Q1. Interestingly, the effects appear to be persistent over time. This provides important evidence from a policy evaluation perspective, as the persistence of prepositioning is in line with the intention of the policy.

In Figure 2, we replicate our analyses on the estimated quarterly coefficients where the continuous measure, LTV, is our dependent variable. Again, the estimated coefficient in 2016:Q4 is close to zero and statistically insignificant. This further confirms that there is no evidence of a violation of the parallel trends assumption. Moreover, the estimated effects remain persistent over time. In other words, the overall LTV composition of new mortgage lending decreased permanently following the publication of the policy statement.

5.3 *The Role of Supervisory Interactions*

Our findings on prepositioning may mask heterogeneity among the group of eligible banks driven by the frequency of their supervisory interactions. It seems plausible that a bank with a higher number of supervisory interactions shall face a lower degree of uncertainty regarding its eligibility status one way or the other. For example, while more interaction may be due to a period of closer scrutiny to address specific concerns raised by supervisors, the affected bank

Figure 1. Prepositioning into Low LTV (<60 percent)




Note: Figure 1 shows the results of a linear regression model estimated using OLS on the mortgage transaction level. The dependent variable is the propensity to grant a low-LTV loan below 60 percent. The coefficients displayed are the time dummy variables interacted with bank eligibility (“below two-thirds” criterion). This figure shows the coefficients for banks that had their SREP in 2019 incl. banks with SREPs in 2017 on a one-/two-year cycle and banks with SREPs in 2018 on a one-year cycle. The first dotted vertical line relates to the quarter of the publication of the CP, 2017:Q1. The second dotted vertical line relates to the quarter of the publication of the PS, 2017:Q4. The third dotted vertical line relates to the last quarter of the year before banks had their SREP assessments, 2018:Q4. The horizontal line illustrates no effect relative to 2016:Q1. Two-way standard errors are clustered on the region and time level. The figure shows 95 percent confidence intervals.

may gain a better understanding of what concrete mitigating steps could be taken to secure eligibility.¹⁸ To further investigate the impact of interaction with supervisors, we utilize unique regulatory data on the bilateral interaction between banks and their supervisors.

¹⁸It is worth pointing out that the two key determinants of ineligibility are supervisory concerns in terms of governance and riskiness of the business model.

Figure 2. Prepositioning into LTV (continuous measure)



Note: Figure 2 shows the results of a linear regression model estimated using OLS on the mortgage transaction level. The dependent variable is the continuous variable, LTV. The coefficients displayed are the time dummy variables interacted with bank eligibility (“below two-thirds” criterion). This figure shows the coefficients for banks that had their SREP in 2019 incl. banks with SREPs in 2017 on a one-/two-year cycle and banks with SREPs in 2018 on a one-year cycle. The first dotted vertical line relates to the quarter of the publication of the CP, 2017:Q1. The second dotted vertical line relates to the quarter of the publication of the PS, 2017:Q4. The third dotted vertical line relates to the last quarter of the year before banks had their SREP assessments, 2018:Q4. The horizontal line illustrates no effect relative to 2016:Q1. Two-way standard errors are clustered on the region and time level. The figure shows 95 percent confidence intervals.

We hypothesize that *eligible* banks with a comparatively high frequency of supervisory interaction are more certain about their eligibility status and, therefore, more likely to be responsive to the publication of the CP, although at the consultation stage banks could have been uncertain about the likelihood of the proposal being confirmed or the likelihood of being eligible. However, as outlined in Section 5.1, the PRA had clearly signaled its firm intention to implement the proposal well in advance of the publication of the CP. Therefore, the only residual source of uncertainty faced by banks

was arguably around their eventual eligibility under the proposed scoping criteria.

In Table 7, we replicate our main analyses for our two dependent variables *LTV (<60 percent)* and *LTV*, in columns 1 and 3, respectively. In the specifications shown in columns 2 and 4, we then add a triple-interacted term combining the time, eligibility, and high supervisory interaction dummy variables (as set out in Section 3.3). We acknowledge that supervisory intensity is not randomly assigned and that banks receiving higher levels of supervision likely differ in relevant ways from those with lower supervision.¹⁹ We address these differences by controlling for time-invariant bank characteristics using bank fixed effects in all columns.

The specification in column 2 shows coefficients of 21 percentage points on the triple interaction term *Eligible* $\times$ *Post (CP)* $\times$ *Supervision (high)* and of -7 percentage points on the double-interacted term *Eligible* $\times$ *Post (CP)*, both at the 5 percent significance level. The specification in column 4, where continuous LTV is our dependent variable, shows similar results. These results suggest that eligible banks with a high frequency of bilateral supervisory interactions responded to the publication of the CP by prepositioning into low-LTV segments, in contrast to the conduct of eligible banks without a high frequency of bilateral interaction.²⁰

Once more, the importance of this result is amplified because the double-interacted term *Post (CP)* $\times$ *Supervision (high)* is of similar magnitude but opposite sign, meaning that noneligible banks with a high frequency of interaction with their supervisors also reacted earlier in response to the publication of the CP, but in an opposite direction to that of eligible banks. Intuitively, this seems sensible. Finally, it is also to be expected that the triple- and double-interacted coefficients with respect to the publication of the PS reverse the estimated effect of the CP publication, given the comparatively later

¹⁹Indeed, while similar in size, banks with high supervisory contact tend to have riskier mortgages, more deposit funding, and lower capital ratios than other banks. These differences in “business risk” align with the rationale that heightened supervisory attention is typically directed toward banks with greater risk exposure.

²⁰We note that the triple interaction term *Eligible* $\times$ *Post (PS)* $\times$ *Supervision (high)* is insignificant though the estimate is almost three times the standard error. This could be because of the way we cluster standard errors.

Table 7. The Heterogeneous
Role of Supervisory Interactions

Dependent Variable	LTV (<60 percent)		LTV	
	(1)	(2)	(3)	(4)
Eligible × Post (CP)	−0.031 (0.015)	−0.072** (0.016)	1.643 (2.586)	1.395** (0.259)
Eligible × Post (PS)	0.127*** (0.008)	0.116** (0.019)	−7.076*** (2.261)	−5.657** (0.904)
Eligible × Post (CP) × Supervision (High)		0.211** (0.027)		−6.174*** (0.555)
Eligible × Post (PS) × Supervision (High)		−0.148* (0.035)		5.251 (1.876)
Post (CP) × Supervision (High)		−0.191*** (0.014)		8.523*** (0.494)
Post (PS) × Supervision (High)		0.223*** (0.022)		−10.076** (1.390)
Bank FE	YES	YES	YES	YES
Region FE	NO	NO	NO	NO
Region × Time FE	YES	YES	YES	YES
Contract Variables	YES	YES	YES	YES
Observations	62,977	62,977	62,977	62,977
Number of Banks	13	13	13	13
Eligible Banks	9	9	9	9
Supervision (High)	5	5	5	5
Mean of Dep. Variable	0.344	0.344	66.433	66.433
R2	0.271	0.272	0.382	0.383
Clustered SE	YES	YES	YES	YES
Method	OLS	OLS	OLS	OLS

Note: This table shows the results of a linear regression model estimated using OLS on the mortgage transaction level. In columns 1 and 2, the dependent variable is the propensity to grant a low-LTV loan below 60 percent. In columns 3 and 4, the dependent variable is the continuous variable, LTV. In all columns, this table shows the results for banks that had their SREP in 2019 (or in 2017 on a one-/two-year cycle or in 2018 on a one-year cycle). Explanatory variables are a time dummy variable as well as its interaction with lender eligibility. In all columns, we control for contract variables and bank fixed effects. In columns 1 and 3, we include region fixed effects. In columns 2 and 4, we include region × time fixed effects and triple interaction terms with supervision (high). Contract variables include fixed rate, incentive period, and refinancing. Standard errors are clustered on the region and time level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 level, respectively.

response of both eligible and noneligible banks with lower degrees of supervisory interaction, as reported in the previous section.

5.4 *The Timing of Policy Implementation via SREP Cycles*

The fact that the prepositioning of banks with a SREP in 2019 persists throughout 2019 may have weakened the ability of the subsequent cohort of eligible banks, those with a SREP in 2020, to do the same. In our main analyses in Table 5, we had shown that the sign, magnitude, and significance of the noninteracted time dummies for the publication of the CP and PS indicated that the prepositioning of eligible banks crowded out, to some extent, noneligible banks within the 2019 SREP cohort. By the same token, the persistence of prepositioning for the 2019 SREP cohort may have prevented eligible banks in the 2020 SREP cohort from doing likewise. This is indeed what we observe when we replicate our main specification comparing eligible and noneligible banks in the 2020 SREP cohort.

To investigate this further, we compare the two groups of eligible banks in the 2019 and 2020 cohorts over the eight consecutive quarters in 2018 and 2019. Compared with Table 5, we replace the control group (banks that had their SREP in 2019 but were ineligible) with banks having their SREP in 2020 but considered eligible. We then compare their prepositioning with banks that had their SREP in 2019 and were considered eligible.

Table 8 shows these results. The coefficient of the interaction term *Eligible* $\times$ *Post (PS)* is positive at the 1 percent significance level across all specifications. This indicates that eligible banks in the 2019 SREP cohort began prepositioning relative to the subsequent cohort of eligible banks, those with a SREP in 2020. In addition, the noninteracted *Post (PS)* time dummy variable has a negative coefficient, albeit of lower magnitude, at the 5 percent significance level in specification (1). These results suggest that eligible banks in the 2020 SREP cohort moved in the opposite direction, although to a lesser extent, supporting the existence of crowding-out effects.

5.5 *Refinements of Our Main Analyses*

In this subsection, we conduct two refinements of our main analyses. First, we present our main results using two alternative thresholds

Table 8. The Impact on Other SREP Cohorts

Dependent Variable	LTV (<60 percent)		LTV	
	(1)	(2)	(3)	(4)
Eligible $\times$ Post (CP)	-0.023 (0.010)	0.004 (0.005)	0.598 (0.315)	0.139 (0.211)
Eligible $\times$ Post (PS)	0.103*** (0.007)	0.101*** (0.007)	-3.999*** (0.341)	-4.262*** (0.315)
Post (CP)	0.025** (0.006)		-0.175 (0.130)	
Post (PS)	-0.078** (0.011)		1.871** (0.355)	
Bank FE	YES	YES	YES	YES
Region FE	YES	NO	YES	NO
Region $\times$ Time FE	NO	YES	NO	YES
Contract Variables	YES	YES	YES	YES
Observations	82,012	82,012	82,012	82,012
Number of Banks	18	18	18	18
SREP 2020 & Eligible Banks	10	10	10	10
SREP 2019 & Eligible Banks	8	8	8	8
Mean of Dep. Variable	0.307	0.307	68.720	68.720
R2	0.261	0.261	0.376	0.376
Clustered SE	YES	YES	YES	YES
Method	OLS	OLS	OLS	OLS

Note: This table shows the results of a linear regression model estimated using OLS on the mortgage transaction level using the subsample. In columns 1 and 2, the dependent variable is the propensity to grant a low-LTV loan (below 60 percent). In columns 3 and 4, the dependent variable is the continuous variable, LTV. In all columns, this table shows the results for banks that had their SREP in 2019 (or in 2017 on a two-year cycle or in 2018 on a one-year cycle) and were ineligible as well as banks that had their SREP in 2020 (or 2018 on a two-year cycle) and were eligible. One bank that had SREPs both in 2019 and in 2020 was left out. Explanatory variables are a time dummy variable as well as its interaction with bank eligibility. In all columns, we control for contract variables and bank fixed effects. In columns 1 and 3, we include region fixed effects. In columns 2 and 4, we include region $\times$ time fixed effects. Contract variables include fixed rate, incentive period, and refinancing. Standard errors are clustered on the region and time level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 level, respectively.

for low LTV in Table 9: below 50 percent in columns 1 and 2, and below 80 percent in columns 3 and 4. When examining LTV below 50 percent, compared to LTV below 60 percent as shown in

**Table 9. The Average Effect of Capital Reliefs on
Prepositioning (alternative dependent variables)**

Dependent Variable	LTV (<50 percent)		LTV (<80 percent)	
	(1)	(2)	(3)	(4)
Eligible × Post (CP)	−0.022 (0.016)	−0.008 (0.011)	−0.086 (0.036)	−0.079 (0.033)
Eligible × Post (PS)	0.077** (0.011)	0.075** (0.010)	0.191*** (0.012)	0.196*** (0.010)
Post (CP)	0.037 (0.024)		0.066 (0.036)	
Post (PS)	−0.068** (0.013)		−0.148** (0.018)	
Bank FE	YES	YES	YES	YES
Region FE	YES	NO	YES	NO
Region × Time FE	NO	YES	NO	YES
Contract Variables	YES	YES	YES	YES
Observations	62,977	62,977	62,977	62,977
Number of Banks	13	13	13	13
Eligible Banks	9	9	9	9
Mean of Dep. Variable	0.257	0.257	0.596	0.596
R2	0.253	0.254	0.307	0.308
Clustered SE	YES	YES	YES	YES
Method	OLS	OLS	OLS	OLS

Note: This table shows the results of a linear regression model estimated using OLS on the mortgage transaction level. In columns 1 and 2, the dependent variable is the propensity to grant a low-LTV loan below 50 percent. In columns 3 and 4, the dependent variable is the propensity to grant a loan with an LTV below 80 percent. In all columns, this table shows the results for banks that had their SREP in 2019. Explanatory variables are a time dummy variable as well as its interaction with bank eligibility. In all columns, we control for contract variables and bank fixed effects. In columns 1 and 3, we include region fixed effects. In columns 2 and 4, we include region × time fixed effects. Contract variables include fixed rate, incentive period, and refinancing. Standard errors are clustered on the region and time level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 level, respectively.

columns 1 and 2 of Table 6, the prepositioning effect is nearly half the size and significant only at the 5 percent significance level. By contrast, columns 3 and 4 of Table 9 show that the effect is much stronger in magnitude and statistically significant at the 1 percent level once we employ the more conservative criterion of LTV below 80 percent.

These results are intuitive. The gaps in risk weights between IRB and SA, which drive the expected benefits in terms of capital relief from the policy implementation, are already material at LTV levels below 80 percent, as shown in Table 1. Therefore, we expect to observe the prepositioning effect across all LTV bands below 80 percent. As the treatment LTV threshold is lowered, the prepositioning effect appears weaker because some observations within LTV levels with an IRB versus SA risk weight difference are reclassified from low LTV to high LTV. For example, as shown in Table 4, 34 percent of mortgage originations are between 50 percent and 80 percent, as opposed to 26 percent at LTV below 50 percent. This distribution implies that eligible banks must rely on relatively higher LTV bands as part of their prepositioning adjustments.

In summary, the results presented in Table 9 confirm that banks that prepositioned did so across the entire spectrum of low LTV bands where a gap in risk weights between SA and IRB exists, rather than just at the very low LTV end.

Second, we replicate our main analyses using an alternative (more conservative) eligibility criterion defining a bank to be eligible if it has an average risk score less than one-half of the maximum risk score. This adjustment reduces the number of eligible banks from nine to three, as shown in Table 10. Looking at the results, both the prepositioning in low LTV and the reduction of the overall LTV level are statistically significant. They are slightly stronger in magnitude compared to the results presented in Table 6 once we control for the interacted region $\times$ time fixed effect (columns 2 and 4). This stronger effect is expected, as the three remaining eligible banks are among the most conservative, meaning they likely faced the least uncertainty regarding their eligibility for capital relief. With greater confidence in the benefits of prepositioning, these banks were more proactive in adjusting their balance sheets, which explains the observed larger effects on lending behavior.

While the number of eligible banks is reduced to three, the fact that the results move in the expected direction—showing a stronger

**Table 10. The Average Effect of Capital Reliefs on
Positioning (alternative eligibility criterion)**

Dependent Variable	LTV (<60%)		LTV	
	(1)	(2)	(3)	(4)
Eligible (Avg. Risk) × Post (CP)	0.008 (0.017)	−0.027 (0.012)	0.856 (0.383)	1.948* (0.604)
Eligible (Avg. Risk) × Post (PS)	0.088** (0.020)	0.138** (0.014)	−5.628** (1.114)	−8.121** (0.883)
Post (CP)	0.034 (0.026)		−1.629 (0.662)	
Post (PS)	−0.057 (0.021)		2.513 (1.026)	
Bank FE	YES	YES	YES	YES
Region FE	YES	NO	YES	NO
Region × Time FE	NO	YES	NO	YES
Contract Variables	YES	YES	YES	YES
Observations	62,977	62,977	62,977	62,977
Number of Banks	13	13	13	13
Eligible Banks	3	3	3	3
Mean of Dep. Variable	0.344	0.344	66.433	66.433
R2	0.280	0.283	0.405	0.407
Clustered SE	YES	YES	YES	YES
Method	OLS	OLS	OLS	OLS
Note: This table shows the results of a linear regression model estimated using OLS on the mortgage transaction level. In columns 1 and 2, the dependent variable is the propensity to grant a low-LTV loan below 60 percent. In columns 3 and 4, the dependent variable is the continuous variable, LTV. In all columns, this table shows the results for banks that had their SREP in 2019. Explanatory variables are a time dummy variable as well as its interaction with lender eligibility. In all columns, we control for contract variables and bank fixed effects. In columns 1 and 3, we include region fixed effects. In columns 2 and 4, we include region × time fixed effects. Contract variables include fixed rate, incentive period, refinancing, and first-time buyer. Standard errors are clustered on the region and time level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 level, respectively.				

response from more conservative banks—supports the robustness of our findings. The prepositioning effect is less strong in the specification with just the regional fixed effect (column 1) but the overall pattern remains consistent with our expectations.

5.6 *Bank Balance Sheet Adjustments*

Building on our previous analysis, we extend our analyses to assess the broader impact of the policy change on bank behavior. A capital relief in one asset class can affect the overall portfolio, as banks adjust their allocation to maximize risk-adjusted returns. Considering this, we examine both the composition of the asset side of banks' balance sheets and changes in total assets.

To that end, we use an updated version of the Bank of England's Historical Banking Regulatory Database (HBRD), which is described by de Ramon, Francis, and Milonas (2017). The HBRD includes detailed bank balance sheet and regulatory data for all PRA-regulated banks in the U.K., offering a more comprehensive view of financial firms, particularly of smaller banks, than other data sets. It provides data on various asset classes, including corporate loans, SME loans, credit card loans, mortgages, and total assets. We use the regression framework outlined in Section 4 to estimate the impact of the PS and CP on banks' adjustments of assets. Our specification includes bank and time fixed effects to control for some unobserved time-invariant heterogeneity. However, we no longer include regional or contract-level fixed effects, as we now use bank balance sheet observations which are at the bank $\times$ time level, whereas previously our data were at the mortgage transaction level.

Table 11 presents our results. The policy resulted in an overall expansion of banks' balance sheets, with eligible banks increasing their total assets by approximately GBP 129 million compared to ineligible banks—a notable rise relative to the sample average of GBP 1.8 billion (column 1). We also observe significant changes in the composition of banks' balance sheets. While corporate loans remained unaffected, eligible banks increased their investments in both SME loans and mortgages. Specifically, eligible banks raised their SME lending by 2.6 percentage points (column 3) and expanded their mortgage portfolio share by 4.2 percentage points

Table 11. The Average Effect of Capital Reliefs on Balance Sheet Adjustments

Dependent Variable	Total Assets (GBP m)	Corporate Loans (% of Total Assets)	SME Loans (% of Total Assets)	Card Loans (% of Total Assets)	Mortgages (% of Total Assets)
	(1)	(2)	(3)	(4)	(5)
Eligible × Post (CP)	120.289 (78.105)	−0.436 (0.411)	0.140 (1.147)	−0.007 (0.008)	3.204 (2.448)
Eligible × Post (PS)	129.444** (53.073)	−0.417 (0.382)	2.636*** (0.958)	−0.016*** (0.006)	4.234*** (1.457)
Bank FE	YES	YES	YES	YES	YES
Time FE	YES	YES	YES	YES	YES
Observations	167	167	167	167	167
Number of Banks	12	12	12	12	12
Eligible Banks	9	9	9	9	9
Mean of Dep. Variable	1,821.289	5.835	2.580	0.039	69.612
R2	0.993	0.981	0.777	0.965	0.882
Robust SE	YES	YES	YES	YES	YES
Method	OLS	OLS	OLS	OLS	OLS
Note: This table shows the results of a linear regression model estimated using OLS on the bank balance sheet level. In all columns, this table shows the results for banks that had their SREP in 2019 as well as banks with SREPs in 2017 on a one-/two-year cycle and banks with SREPs in 2018 on a one-year cycle. In all columns, we control for year-quarter and bank fixed effects. In all columns, one bank is left out due to data quality concerns. In column 1, we show total assets in nominal terms. In columns 2–5, we standardize by total assets. Standard errors are heteroskedasticity-robust. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 level, respectively.					

(column 5), both relative to ineligible banks. Conversely, we observe a slight (0.02 percentage point) reduction in credit card loan exposure (column 4).

Our results reveal banks' dual strategy in response to the capital relief. They increased investments in riskier, higher-return SME loans while simultaneously expanding their holdings in safer, lower-return mortgages, alongside an overall increase in total assets. At the same time, they avoided the highest-risk asset class—credit card loans. These findings reinforce our analysis of the policy change's significant impact on both the composition and size of banks' balance sheets.

6. Application: Estimating the Shadow Cost of Capital for Small Banks

This section discusses how our estimates of the adjustment in the portfolio allocation ahead of the SREP can be used to quantify a lower bound for the shadow cost of small banks' regulatory equity capital requirements. This is because the portfolio adjustment incurs costs, as it entails a departure from what must have been the optimal portfolio allocation under the existing regulatory treatment.²¹ Thus, the cost of prepositioning is expected to be lower than the benefits from the expected capital relief that could be accrued by the bank because of the following supervisory review. Regarding the cost of prepositioning, the determining factor is the cost of equity capital, which is typically unobservable for small and unsophisticated banks targeted by this policy, as they usually do not raise equity capital externally. In this section, we propose a way to gauge the minimum (unobservable) cost of equity capital faced by these opaque banks using the (observable) cost of prepositioning, which we estimate.

To estimate this, we assume that during prepositioning eligible banks increase their lending by raising more liabilities (i.e., typically in the form of retail deposits) because lending at lower LTV requires less regulatory capital. This assumption is supported by our evidence on banks' balance sheet expansion presented in the previous section.

²¹We are implicitly assuming that the current optimal portfolio allocation is constant over the prepositioning period.

Moreover, we assume that banks' changes in their portfolio allocations are restricted to the mortgage book. However, this assumption is conservative and leads to a lower bound estimate of the cost of capital considering the finding presented in the previous section. We showed that eligible banks increased SME lending, presumably to mitigate the reduction in profitability caused by prepositioning. We also implicitly assume that banks can raise additional liabilities to fund the lending expansion at the same marginal cost. However, we expect small banks to offer higher rates to raise deposits, which suggests that the cost of funding the lending expansion thereof is increasing. Indeed, the finding that eligible banks increased SME lending, which is subject to high risk weights, is consistent with the view that banks find it costly to raise extra funding to support lending expansion. As these latter assumptions tend to offset each other, we believe that our simplification is defensible. Accordingly, under the SA, the risk weight for low-LTV residential mortgages is slightly lower than the one at high LTV: $RW_{SA}^H > RW_{SA}^L$. This means that for any unit of lending at high LTV, they can invest $\frac{RW_{SA}^H}{RW_{SA}^L} > 1$ units at low LTV.

From the previous paragraph it follows that an eligible bank incurs the cost of prepositioning in the initial period between the announcement of the policy and its SREP and it is driven by the opportunity cost of investing in a high-risk asset as compared to a low-risk asset. The opportunity cost is due to a lower return for low-LTV mortgages (i.e., $r_L < r_H$).²² However, this is partly offset by a slightly higher nominal amount of the loan, as the risk weight for low-LTV mortgages is lower (i.e., $RW_{SA}^L < RW_{SA}^H$). Accordingly, the cost of prepositioning can be formulated as follows:

$$\text{Cost of prepositioning} = \Delta\theta_l \left(r_H * \frac{RW_{SA}^L}{RW_{SA}^H} - r_L \right),$$

where $\Delta\theta_l$ is the degree of prepositioning (i.e., the change in the share of lending to low-risk assets), r_H is the return on high-risk

²²In this respect, we implicitly assume that the reduction in interest rate is not at least partly compensated by correspondingly lower provisions for expected losses. We believe that this is a defensible assumption given that, as shown in Table 3, the average LTV for eligible banks before prepositioning was already lower than 60 percent.

assets, and r_L is the return on low-risk assets. We assume that banks that preposition by expanding their exposure at low LTV maintain that allocation until the following supervisory review (i.e., in line with the aim of the policy).

Eligible banks benefit from a capital relief at the time of the supervisory review which happens after the announcement of the policy. For capital budgeting purposes, we assume that these banks can reinvest the extra amount of equity resources at the corresponding cost of equity capital, r_E , until the following supervisory review. Therefore, both the cost of prepositioning and the expected benefits from it are subject to the same time frame corresponding to the three-year SREP cycle. This simplification allows us to ignore discounting on both sides. Therefore, we arrive at the following inequality:²³

$$\Delta\theta_l \left(r_H * \frac{RW_{SA}^L}{RW_{SA}^H} - r_L \right) \leq E[Capital\ relief] * r_E.$$

6.1 Numerical Example

In this subsection, we provide a numerical example of the cost of capital calculation. To derive both the ratio $\frac{RW_{SA}^L}{RW_{SA}^H}$ and the expected capital relief, we rely on the table published by the regulator comparing risk weights between SA and IRB, as shown in Table 1.

RW_{SA}^L is equal to 35 percent. To calculate RW_{SA}^H , we simply take the average of the SA risk weights for the four LTV bands above 60 percent, which results in an average risk weight of 37.25 percent and a ratio of 0.94.²⁴

Regarding the expected capital relief, we rely on an indication provided by the regulator in its 2018 Annual Competition Report that “on average, we anticipate these [banks]’ minimum capital requirements will reduce by 10%–15%” (PRA 2018b). Assuming a total capital requirement of 10 percent and a risk-weighted asset

²³This approach rests on the underlying simplifying assumption that banks are monoline mortgage banks, i.e., they only invest in mortgage loans.

²⁴We ignore the residual band with $LTV \geq 100$ percent due to a low number of mortgages in the sample period.

density of 36.5 percent,²⁵ the expected capital relief is estimated to be 0.36 to 0.55 percent of total assets, which is directly comparable with the cost of prepositioning, also expressed as a percentage figure out of total assets.

For illustrative purposes, we assume that “Bank A” has the following return on assets: $r_H = 3.6$ percent and $r_L = 3.1$ percent. Its expected capital relief is 0.50 percent and its degree of prepositioning is 12 percent. This would yield a lower bound of the cost of capital r_E of 6.8 percent:

$$r_E \geq \frac{12\% * \left(3.6\% * \frac{35\%}{37.25\%} - 3.1\% \right)}{0.50\%} = 6.8\%.$$

6.2 *Estimates of Bank-Specific Lower Bound of Capital Cost*

We then estimate the bank-specific cost of capital for the subset of banks with SREPs in 2019 that were eligible. To do so, we estimate multiple regressions by comparing the degree of prepositioning bank-by-bank relative to all banks in the control group. We use the 60 percent LTV threshold because this LTV band combines both a wide IRB versus SA gap (i.e., expected benefits) and the largest differential in interest rates (i.e., profit sacrifice). We observe that the degree of prepositioning ranges from 10 percent to 21 percent, with an average of 15 percent.²⁶ Similarly, we obtain bank-specific estimates of capital reliefs from internal policy impact assessments ahead of the publication of the policy statement. They range from 0 percent of total assets to 0.9 percent of total assets. To proxy the return on high-LTV and low-LTV mortgages, we use the average rates of two-year fixed mortgages in the year 2016 for mortgages below 60 percent LTV and those with at least 60 percent LTV. We find that these shadow costs of capital differ across banks, ranging from 7 percent to 16 percent, with an average of 11 percent and a median of 11 percent.²⁷ It is worth noting that this estimate is

²⁵This figure is the average of the SA risk weights across the six LTV bands as shown in Table 1. This approach is in line with the assumption that the relevant banks are monoline mortgage banks.

²⁶We exclude two outliers with no prepositioning at all.

²⁷We also conduct several sensitivity analyses with the average cost of 13 percent, calculated by equally weighting the mortgage rates charged by each bank;

only marginally above the estimates for the cost of capital for large banks.²⁸

7. Conclusions

We study a salient change in policy that creates an incentive for banks to change their asset composition to benefit from a reduction of their capital requirements. We examine whether banks react to anticipated reliefs of capital requirements and to which of the announcements in the policy life cycle they react.

Using a novel data set of the population of all residential mortgages and supervisory risk scores, we demonstrate that banks likely to receive capital relief increased their benefit by investing in low-risk assets (i.e., low-LTV mortgages) following the publication of the policy. In addition, using unique data on the frequency of bank-supervisor interaction, we show that likely eligible banks with a higher number of supervisory interactions reacted earlier, in response to the publication of the consultation paper. We interpret this as indicative evidence that the bilateral communication between regulated banks and their supervisors allowed the former to better gauge whether they were likely to be in scope of, and thus benefit from, the policy in question.

From a policy evaluation perspective, we find not only that the initial response of eligible banks, increasing their exposure to lower-risk assets, aligns with the direction intended by the policy, but also that the response persists beyond its implementation. Such persistence rules out opportunistic gaming that would have been problematic from a prudential perspective. However, we detect the opposite effect for ineligible banks, namely reducing their exposure to lower-risk assets. Similarly, we find that eligible banks with more time to change their asset composition did not invest in low-risk assets once those in the previous cohort had already repositioned.

this reduces to 5.5 percent when considering five-year fixed mortgages in the year 2016.

²⁸For example, the latest survey by the European Banking Authority estimates banks' cost of equity to be in the range of 8 to 10 percent; see EBA (2019).

Table A.1. Examples of Wording in Supervisory Review Letters

Low Risk	• <i>It would be helpful in future iterations of its ICAAP if the bank explained how this figure, and the calibration of the trigger for action, were arrived at.</i> • <i>We encourage the bank to consider how the risk function will develop to provide proactive challenge and take full ownership of the bank's risk framework.</i> • <i>The bank's credit risk appetite should be better articulated.</i>
Medium Risk	• <i>Further evolution of MI is however required, with appropriate commentary to facilitate Board discussion.</i> • <i>Progress has been made on the development of the Risk Management Framework (RMF). (...) The bank recognized that while the new RMF is materially embedded, it cannot be fully so until all the new executive hires are in place and the ongoing cultural changes embedded.</i> • <i>Whilst the continued commitment of the NEDs has been critical in achieving progress, we believe that this arrangement is likely to be unsustainable in the medium to longer term and that it may – if continued – compromise their independence.</i> • <i>We ask that the scope of their remit is widened to include a review of the bank's pricing model, in addition to other financial models too, in line with good practice and governance.</i>
High Risk	• <i>The focus of our supervisory strategy for the next 12 months will therefore continue to be the bank's development of a long-term strategy to accrete capital resources and the ongoing management of the risks of its loan book.</i> • <i>We would therefore ask you to review your approach to hedging FX risk as a matter of urgency and provide us with the output of your analysis by the end of October.</i>

References

- Andrade, P., G. Gaballo, E. Mengus, and B. Mojon. 2019. "Forward Guidance and Heterogeneous Beliefs." *American Economic Journal: Macroeconomics* 11 (3): 1–29.
- Bank of England. 2023. "Mind the (Smaller) Gap? Implications of the Narrowing Gap between Modelled and Standardised Residential Mortgage Risk Weights." *Bank Overground*, February 27. bankofengland.co.uk/bank-overground/2023/implications-of-the-gap-between-modelled-and-standardised-residential-mortgage-risk-weights.
- Basel Committee on Banking Supervision (BCBS). 2001. "Pillar 2 (Supervisory Review Process)." Consultative Document, Bank for International Settlements. <https://www.bis.org/publ/bcbsca08.pdf>.
- Benetton, M., P. Eckley, N. Garbarino, L. Kirwin, and G. Latsi. 2021. "Capital Requirements and Mortgage Pricing: Evidence from Basel II." *Journal of Financial Intermediation* 48: Article 100883.
- Berrospide, J. M., and R. M. Edge. 2024. "Bank Capital Buffers and Lending, Firm Financing and Spending: What Can We Learn from Five Years of Stress Test Results?" *Journal of Financial Intermediation* 57 (C): Article 101061.
- Bholat, D., J. Brookes, C. Cai, K. Grundy, and J. Lund. 2017. "Sending Firm Messages: Text Mining Letters from PRA Supervisors to Banks and Building Societies They Regulate." Staff Working Paper No. 688, Bank of England.
- Bruno, B., E. Onali, and K. Schaeck. 2018. "Market Reaction to Bank Liquidity Regulation." *Journal of Financial and Quantitative Analysis* 53 (2): 899–935.
- Buch, C., M. Buchholz, and L. Tonzer. 2015. "Uncertainty, Bank Lending, and Bank-Level Heterogeneity." *IMF Economic Review* 63 (4): 919–54.
- Cloyne, J., K. Huber, E. Ilzetzki, and H. Kleven. 2019. "The Effect of House Prices on Household Borrowing: A New Approach." *American Economic Review* 109 (6): 2104–36.
- de Ramon, S. J., W. B. Francis, and K. Milonas. 2017. "An Overview of the UK Banking Sector since the Basel Accord: Insights from a

- New Regulatory Database.” Staff Working Paper No. 652, Bank of England.
- Eckley, P., W. Francis, and A. Mathur. 2021. “In the Dangerzone! Regulatory Uncertainty and Voluntary Bank Capital Surpluses.” Mimeo.
- European Banking Authority (EBA). 2014. “Guidelines on Common Procedures and Methodologies for the Supervisory Review and Evaluation Process (SREP).” EBA/GL/2014/13. [https://eba.europa.eu/sites/default/documents/files/documents/10180/935249/4b842c7e-3294-4947-94cd-ad7f94405d66/EBA-GL-2014-13%20\(Guidelines%20on%20SREP%20methodologies%20and%20processes\).pdf](https://eba.europa.eu/sites/default/documents/files/documents/10180/935249/4b842c7e-3294-4947-94cd-ad7f94405d66/EBA-GL-2014-13%20(Guidelines%20on%20SREP%20methodologies%20and%20processes).pdf).
- . 2019. “Risk Assessment Questionnaire – Summary of the Results.” Report, June. <https://eba.europa.eu/sites/default/documents/files/documents/10180/2854739/916f8c4b-7099-4aba-ac1f-882cfd4c3583/RAQ%20Booklet%20Spring%202019.pdf>.
- Financial Conduct Authority (FCA). 2018. “Strategic Review of Retail Banking Business Models.” <https://www.fca.org.uk/publication/multi-firm-reviews/strategic-review-retail-banking-business-models-final-report.pdf>.
- Fraisse, H., M. Lé, and D. Thesmar. 2020. “The Real Effects of Bank Capital Requirements.” *Management Science* 66 (1): 5–23.
- Gropp, R., H. Hakenes, and I. Schnabel. 2011. “Competition, Risk-Shifting, and Public Bail-out Policies.” *Review of Financial Studies* 24 (6): 2084–120.
- Gropp, R., T. Mosk, S. Ongena, and C. Wix. 2019. “Bank Response to Higher Capital Requirements: Evidence from a Quasi-Natural Experiment.” *Review of Financial Studies* 32 (1): 266–99.
- Imbens, G. W., and D. B. Rubin. 2015. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press.
- Imbens, G. W., and J. M. Wooldridge. 2009. “Recent Developments in the Econometrics of Program Evaluation.” *Journal of Economic Literature* 47 (1): 5–86.
- Irani, M. I., R. Iyer, R. R. Meisenzahl, and J.-L. Peydró. 2021. “The Rise of Shadow Banking: Evidence from Capital Regulation.” *Review of Financial Studies* 34 (5): 2181–235.

- Jokipii, T., and A. Milne. 2008. "The Cyclical Behaviour of European Bank Capital Buffers." *Journal of Banking & Finance* 32 (8): 1440–51.
- Kara, G., and Y. Yook. 2019. "Policy Uncertainty and Bank Mortgage Credit." BIS Working Paper No. 820, October.
- Kisin, R., and A. Manela. 2016. "The Shadow Cost of Bank Capital Requirements." *Review of Financial Studies* 29 (7): 1780–820.
- McKay, A., E. Nakamura, and J. Steinsson. 2016. "The Power of Forward Guidance Revisited." *American Economic Review* 106 (10): 3133–58.
- Prudential Regulation Authority (PRA). 2017. "Refining the PRA's Pillar 2A Capital Framework." Consultation Paper No CP3/17. <https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/consultation-paper/2017/cp317>.
- . 2018a. "The Prudential Regulation Authority's Approach to Banking Supervision." Report, Bank of England (chapter 3). <https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/approach/banking-approach-2018.pdf>.
- . 2018b. "Annual Report 1 March 2017–28 February 2018." Technical Report, June. <https://www.bankofengland.co.uk/-/media/boe/files/annual-report/2018/pr-2018.pdf?la=en&hash=929BC2A486101460E1A371FF96F0DC73B424BF0F>.
- . 2019. "The Internal Capital Adequacy Assessment Process (ICAAP) and the Supervisory Review and Evaluation Process (SREP)." Supervisory Statement No. SS31/15. <https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/supervisory-statement/2019/ss3115-update-march-2019>.
- . 2020. "The PRA's Methodologies for Setting Pillar 2 Capital – Statement of Policy." <https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/statement-of-policy/2020/the-pras-methodologies-for-setting-pillar-2a-capital-update-february-2020.pdf?la=en&hash=4EB02F435F5BF46507222B09F0DE271333A5409E>.
- . 2023. "The Prudential Regulation Authority's Approach to Banking Supervision." Report, Bank of England. <https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/approach/banking-approach-2023.pdf>.

- Repullo, R., and J. Suarez. 2004. "Loan Pricing under Basel Capital Requirements." *Journal of Financial Intermediation* 13 (4): 496–521.
- Rime, B. 2005. "Will Basel II Lead to a Specialization of Unsophisticated Banks on High-Risk Borrowers?" *International Finance* 8 (1): 29–55.
- Rodnyansky, A., and O. M. Darmouni. 2017. "The Effects of Quantitative Easing on Bank Lending Behavior." *Review of Financial Studies* 30 (11): 3858–87.
- Ruthenberg, D., and Y. Landskroner. 2008. "Loan Pricing under Basel II in an Imperfectly Competitive Banking Market." *Journal of Banking and Finance* 32 (12): 2725–33.
- Schäfer, A., I. Schnabel, and B. W. di Mauro. 2016. "Bail-in Expectations for European Banks: Actions Speak Louder Than Words." CEPR Discussion Paper No. 11061. <https://cepr.org/publications/dp11061>.
- Soto, P. E. 2021. "Breaking the Word Bank: Measurement and Effects of Bank Level Uncertainty." *Journal of Financial Services Research* 59 (1–2): 1–45.
- Stolz, S., and M. Wedow. 2011. "Banks' Regulatory Capital Buffer and the Business Cycle: Evidence for Germany." *Journal of Financial Stability* 7 (2): 98–110.
- Valencia, F. 2016. "Bank Capital and Uncertainty." *Journal of Banking & Finance* 69 (S1–S9): 51–59.
- Woods, S. 2016. "The Revolution Is Over. Long Live the Revolution!" Speech at the City Banquet, Mansion House, London, October 28. <https://www.bankofengland.co.uk/speech/2016/the-revolution-is-over-long-live-the-revolution>.

Understanding the Joint Dynamics of Inflation and Wage Growth in the Euro Area*

Vahagn Galstyan
Central Bank of Ireland

This paper presents an empirical framework and analysis of the interactions among inflation, wages, employment, and output in the euro area. Results identify price shocks and demand shocks as the primary exogenous factors explaining historical variance. The wage gap emerges as a key determinant of wage dynamics in the aftermath of a price shock. In contrast, the output gap becomes dominant following demand shocks. The real wage gap acts as a corrective mechanism, ensuring that prices and wages in particular align with the broader economic landscape. Forecasts for the period starting 2023:Q4 emphasize the enduring significance of the real wage gap, projecting its ongoing impact on nominal wages in tight labor markets. As for inflation expectations, the estimates emphasize their stickiness. In this context, the significant and persistent price shock that has occurred suggests a gradual decline in expectations, potentially leading to an extended period of elevated inflation.

JEL Codes: E00, E12, E30, E31, E32, E37.

1. Introduction

The resurgence of inflation in 2022, following a prolonged period of subdued price dynamics, represents one of the most critical macro-economic developments of the post-pandemic era. The euro area, in particular, witnessed an annualized price growth of 10.6 percent in October 2022, a significant jump from near-zero levels recorded just two years prior. These pressures erode purchasing power, heighten

*I thank Philip Lane, Reamon Lydon, Enrico Sette, and two anonymous referees for invaluable comments and suggestions. Aurelio Nocera provided excellent research assistance. The views expressed in this paper are personal and do not represent the views of the Central Bank of Ireland. Author email: vahagn.galstyan@centralbank.ie.

economic uncertainty, and threaten overall macroeconomic stability. Given the considerable costs of high inflation, understanding its underlying dynamics through empirical analyses becomes imperative. This paper contributes to the ongoing debate on inflation dynamics by analyzing the drivers of inflation and wage growth in the euro area.

Drawing on New Keynesian theory, this paper develops an empirical framework that models demand and supply in both the goods and labor markets. The primary focus is on medium-run cyclical dynamics; however, capturing these requires accounting for both short-term fluctuations and long-term trends. For simplicity, the empirical framework adopts the specification set out in Galí (2015), which comprises four equations. The first equation links inflation to inflation expectations, the deviation of medium-run output from its long-run trend, and a term representing error correction through the real wage gap. The second describes nominal wage growth as a function of inflation expectations, labor market slack, and the real wage gap. The third equation describes aggregate demand as a function of both current and expected inflation. The fourth is a version of Okun's law. This structure combines theoretical rigor with empirical flexibility, making it well-suited for analyzing macroeconomic interactions and consistent with established modeling approaches.¹

This approach shares conceptual similarities with the dynamic empirical model of U.S. inflation developed by Bernanke and Blanchard (2023), which links prices, wages, and inflation expectations using a partial equilibrium framework. They use the vacancy-to-unemployment ratio as a measure of economic slack to assess the effects of both product market and labor market shocks during the pandemic, attributing the 2021 inflation spike primarily to commodity-driven price shocks.² However, the approach taken in this paper differs from their methodology in several important respects. In particular, instead of survey-based expectations, inflation expectations are modeled using only past data. Galstyan (2021)

¹The model builds on the New Keynesian Phillips curve (Woodford 2003), with related frameworks by Calvo (1983), Mankiw and Reis (2002), Gertler and Leahy (2008), and Gabaix (2020). See Eser et al. (2020) on the role of the Phillips curve in monetary policy at the ECB.

²Refer to the appendix for details on Bernanke and Blanchard (2023).

provides empirical support for this backward-looking specification, showing that lagged headline inflation explains a large share of variation in short-term expectations. While this modeling approach may reduce forecasting accuracy during periods of high inflation, it remains a pragmatic choice given the relatively stable inflation environment over much of the sample period.

Furthermore, this paper treats domestic economic slack (both the output and employment gaps) as endogenous, allowing feedback effects between macroeconomic conditions and inflation. By explicitly distinguishing the output gap from the employment gap (linked via a modified version of Okun's law), the model clarifies distinct channels: labor market tightness directly influences wage dynamics, whereas aggregate demand affects labor demand. This distinction enables a more detailed analysis of how shocks originating in goods and labor markets influence inflation dynamics. Additionally, to preserve tractability, the model adopts a parsimonious lag structure that prioritizes richer interactions over extended historical dependencies. Finally, although it uses equation structures similar to those in the work of Bernanke and Blanchard, it further filters out volatility from latent medium-run and long-run variables. This temporal separation helps isolate the "underlying" inflation trend and sheds light on the persistent drivers of inflation.

These modeling choices also position the paper well within several strands in the inflation literature. For instance, Lorenzoni and Werning (2023) develop a model in which inflationary spiral can emerge from disagreements between workers and firms over real wage levels, with the disagreements tied to the magnitude of the output gap. Their analysis shows that even if real wages move in response to shocks, this does not necessarily confirm the presence of a wage-price spiral. The current paper contributes to this literature by explicitly modeling the real wage gap as a key adjustment mechanism. This gap serves as an error-correction term, allowing the model to capture delayed wage responses following relative price shocks—an especially relevant feature in the euro-area context.

Looking at the role of inflation expectations in shaping inflation, Werning (2022) argues that the influence of expectations depends on their time horizon: near-term expectations have a much stronger effect on inflation than longer-term ones. He also shows that the pass-through varies across pricing models and is not necessarily close

to one. These insights are reflected in the empirical findings of this paper, which suggest that medium-run expectations play an important but moderated role, with estimated pass-through coefficients well below one.

Another area of research examines nonlinearities in the relationship between labor market tightness and wage or price dynamics. Benigno and Eggertsson (2023) model nonlinear effects of labor market conditions on inflation, and attribute much of the 2020s U.S. inflation to labor market tightness. Hahn (2020) and Harding, Lindé, and Trabandt (2023) similarly highlight asymmetric Phillips curve responses, especially in periods of low slack or high inflation. De Santis and Tornese (2023) show that energy shocks can transmit more forcefully to consumer prices under high inflation regimes, suggesting further state dependence. In contrast, Bernanke and Blanchard (2023) find that a linear model provides a good approximation of U.S. inflation dynamics. Similarly, recent evidence by Lo Bello and Viviano (2024) finds little support for pronounced nonlinearities in the Phillips curve for the euro area, providing justification for a linear empirical specification adopted in this paper.

Finally, a set of studies focuses specifically on the structural determinants of inflation in the euro area. Lane (2022) delves into the challenges of predicting medium-run inflation trajectories, especially against the backdrop of rising inflation combined with continual external shocks. He singles out the mid-2021 relative price disturbances as major influencers of the inflationary surge, further exacerbated by wage and price rigidities which led to a noticeable uptick in the inflation rate. Based on Shapiro (2022), Gonçalves and Koester (2022) decompose inflation into demand-driven and supply-driven components, showing that inflation during the early stages of the pandemic recovery was dominated by supply shocks, including energy prices and labor shortages, but gradually shifted toward demand factors, especially in services. Bobeica, Ciccarelli, and Vansteenkiste (2019) examine the price-wage relationship in the euro area, finding that the strength of pass-through depends on the nature of the shock and the inflation regime. Hansen, Toscani, and Zhou (2023) highlight the role of import prices and profit margins in explaining recent euro-area inflation.

Summarizing the findings, the shock decomposition reveals that price shocks and shocks to aggregate demand are the principal

exogenous factors explaining historical variance. This is further enriched by the factor decomposition. Between 2019:Q1 and mid-2021, the euro area experienced shifts in inflation dynamics, with the COVID-19 pandemic creating significant disruptions in 2020. While the labor market was robust pre-pandemic, supporting wage growth, the onset of COVID-19 suppressed wages and output. By 2023:Q3, a synergy of price shocks, rising inflation expectations, a rebounding output gap, and labor market recovery pushed inflation and wage growth beyond their pre-pandemic levels.

The impulse response analysis shows that a price shock causes an immediate increase in inflation, which in turn creates a negative real wage gap and stimulates employment. However, elevated inflation contracts aggregate demand. As the effects of the initial shock diminish, inflation remains persistent due to the rise in inflation expectations. On the other hand, a demand shock prompts a positive output gap, subsequently raising inflation and labor demand. Both shocks exhibit a gradual return of variables to their steady-state values, with distinct pathways of adjustment and varying factor contributions. Importantly, the findings emphasize the role of the wage gap in shaping wage dynamics after a price shock, while the output gap is dominant in reactions to a demand shock.

Using the empirical framework for projections of key economic variables from 2023:Q3 onward, the output gap at the beginning of the forecast is approximately 2.2 percent, signaling overheating. The employment gap starts at around 2.0 percent, indicating a tightness in the labor market. This is largely attributed to the positive output gap and the negative real wage gap. The real wage gap is projected to gradually revert to its steady-state value, starting from negative 0.3 percent at the beginning of the forecast horizon.

In relation to inflation and wages, the 2023 year-on-year average inflation rate (GDP deflator) is expected to be 5.2 percent. It is influenced primarily by inflation expectations, followed by the persistence of the price shock and a positive output gap. The estimates emphasize the significant stickiness of inflation expectations. Given the pronounced and enduring price shock observed, this suggests a gradual decline in expectations, potentially leading to an extended period of elevated inflation, with mean predictions for 2024, 2025, and 2026 at 3.8, 2.9, and 2.5 percent, respectively. Meanwhile, wage growth for 2023 is projected to be 5.3 percent. For 2024, 2025, and 2026,

forecasts are 4.3, 3.5, and 3.0 percent, respectively. The dynamics of adjustment are driven largely by the corrective effects of the real wage gap and labor market tightness, with inflationary expectations also playing a role.

The paper is organized as follows. Section 2 introduces the data. Section 3 provides details on the empirical specification. Section 4 discusses the selection of priors. Section 5 presents the results for estimated parameters and extracted medium-run trends. Section 6 delves into a discussion on various decompositions. Section 7 covers the analysis of impulse responses. Section 8 offers out-of-sample projections for the main variables, and Section 9 concludes.

2. Data

The empirical analysis requires quarterly data on inflation, and growth rates of wages, output, and employment. All these variables are derived from data sourced from Eurostat.³

Both the CPI and the GDP deflator present unique advantages. The CPI tracks the prices consumers pay for a specific basket of goods and services, while the GDP deflator offers a broader perspective, capturing a wide range of domestically produced goods. Due to this expansive view, the GDP deflator is adept at reflecting the overall price movements of domestic goods. Notable studies, such as those by Ireland (2004) and Smets and Wouters (2007), used the GDP deflator to explore inflation dynamics over business cycles. As a result, the series for the GDP deflator index is used to construct the measure of inflation.⁴

Wages are measured using per-employee compensation, calculated from the ratio of employee compensation to the number of employees.⁵ Data for total employment are used to calculate the growth rate of employment.⁶ Chain-linked GDP in constant 2015 prices is used to measure the growth rate of output.⁷

³The data were downloaded on 03/07/2024.

⁴Series code: PD15_EUR,SCA,B1GQ,EA19.

⁵Series codes: CP_MEUR,SCA,D1,EA19 and THS.PER,SCA,SAL_DC,EA19.

⁶Series code: THS.PER,SCA,EMP_DC,EA19.

⁷Series code: CLV15_MEUR,SCA,B1GQ,EA19.


Table 1. Summary Statistics

	(1) π_t^p	(2) π_t^w	(3) Δn_t	(4) Δy_t
Sample Mean	0.461	0.561	0.214	0.362
Sample St. Dev.	(0.364)	(0.388)	(0.319)	(0.730)
OLS St. Err.	(0.034)***	(0.036)***	(0.030)***	(0.068)***
Observations	114	114	114	114
Note: The first two rows of the table provide in-sample means and standard deviations for the four observable macroeconomic variables. The third row presents the OLS standard error from a regression of the variable in question on a constant. Robust standard errors are shown in parentheses. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels, respectively.				

The data set, in terms of levels of the variables, spans from the first quarter of 1995 to the third quarter of 2023. This time frame encompasses the COVID-19 era. In order to maintain this timeline while minimizing the impact of extreme fluctuations during this period on the outcomes, a smoothing technique is applied to the variables. Specifically, for all variables measured in logarithmic differences, the Hodrick-Prescott filter with a smoothing parameter of one is applied between 2020:Q1 and 2020:Q3.⁸ These smoothed growth rates are then used to adjust the level observations, ensuring that the smoothing process does not alter the levels of the variables after the smoothing period. The differences in the logarithmic series form the final data covering the period from 1995:Q2 to 2023:Q3.

Table 1 presents summary statistics for the four key variables: inflation (π_t^p), nominal wage growth (π_t^w), employment growth (Δn_t), and real output growth (Δy_t). The sample means of these variables stand at 0.46, 0.56, 0.21, and 0.36, respectively, paired with corresponding sample standard deviations of 0.36, 0.39, 0.32, and 0.73. The high standard deviations in growth rates of output and

⁸Using a smoothing parameter equal to one results in the least possible smoothing. In this scenario, the fluctuations observed during the COVID period are reduced to a magnitude somewhat similar to those of the Global Financial Crisis. Thus, while the series is smoothed, it is not smoothed sufficiently to entirely eliminate the effects of the COVID period.

Figure 1. Time Horizon

Note: The figure provides a diagrammatic representation of the time horizon for a generic variable x_t . This horizon categorizes economic time frames into the steady state, long run, medium run, and short run.

employment are predominantly attributable to the substantial contractions during the Global Financial Crisis and the euro-area debt crisis. Toward the end of the sample period, all variables display pronounced volatility, largely due to the impact of the COVID-19 pandemic.

3. An Empirical Framework

3.1 On Temporal Dimension

The system comprises both observable and unobservable variables. Each variable exists across various time dimensions. Figure 1 offers a diagrammatic representation of the time horizon for a generic variable x_t . This horizon categorizes economic time frames into steady state, long run, medium run, and short run.

From the bottom of the figure, the steady state can be considered a specific subset of the long run, but devoid of perturbations. In such a setting, growth rates for all variables remain constant. The value of a variable in steady state is denoted with a bar over the variable and lacks a time subscript. With inflation being a policy

variable, and in alignment with the European Central Bank (ECB) objective, its value in steady state is set at 0.5 percent (2 percent annualized). Furthermore, in line with the neoclassical growth theory, the methodology presumes that, in steady state, both real wages and output per employee grow at an identical rate: $\Delta\bar{y} - \Delta\bar{n} = \Delta\bar{\omega}$.

In the long run, nominal frictions are absent, and macroeconomic variables remain unaffected by shocks from the short and medium run. This is consistent with the classical dichotomy, suggesting that monetary policy does not exert its impact on real economic factors in this time frame. Here, the long-run value of x_t is denoted by $\bar{x}_t$. In the long run, it is assumed that output, employment, and real wages follow a random walk with drift, where the drift terms across these three variables are constrained in a manner that, in the absence of shocks, the growth rates of these variables revert back to their steady-state levels. Furthermore, it is assumed that a long-run relationship exists between output, employment, and real wages, which can deviate from linear trends. In particular, for $\bar{x}_t \in (\bar{y}_t, \bar{\omega}_t, \bar{n}_t)$ it is assumed that $\Delta\bar{x}_t = \mu^x + \kappa^x \chi_t$, where $\chi_t = \rho_\chi \chi_{t-1} + e_t^\chi$. In the long run, inflation remains fixed at $\bar{\pi}_t^p = \bar{\pi}^p = 0.5$. Long-run inflation expectations also align with the 0.5 benchmark, underscoring the directive role of monetary policy. The growth rate of nominal wages at this frequency is pinned down by the long-run inflation rate and the long-run growth rate of real wages: $\bar{\pi}_t^w = \Delta\bar{\omega}_t + \bar{\pi}^p$.

Central to this paper is the medium run, where nominal frictions come into play and the classical dichotomy breaks. In the medium run, nominal shocks exert an impact on both nominal and real variables. As medium-run shocks exclude high-frequency disturbances, the medium-run trend provides a clearer representation of the dynamic trajectory of variables. Specifically, when discussing inflation, the concept of “underlying inflation” emerges as a medium-run phenomenon, stripping away short-run volatility to uncover foundational, persistent price movements. Consequently, for monetary policy, the future interest rate path is shaped by the medium-run trajectory of inflation. For the purpose of modeling, the medium-run trend of variable x_t is denoted by $\tilde{x}_t$. In this time frame, variables fluctuate around their long-run trend. For output, employment, and real wages, these fluctuations are denoted by a hat, and are referred to as the output gap, employment gap, and real wage gap. Mathematically, these gaps are defined by the

following identities: $\hat{y}_t - \hat{y}_{t-1} = \Delta \tilde{y}_t - \Delta \bar{y}_t$, $\hat{n}_t - \hat{n}_{t-1} = \Delta \tilde{n}_t - \Delta \bar{n}_t$, and $\hat{\omega}_t - \hat{\omega}_{t-1} = \tilde{\pi}_t^w - \tilde{\pi}_t^p - \Delta \bar{\omega}_t$. The rest of the system in the medium run is described in the following subsection.

In the short run, shocks can create significant volatility without affecting the medium-run trend of the variables. Given the limited relevance of this time frame to the current paper, short-run fluctuations are modeled as statistical noise and represented by $\varepsilon_t^x = x_t - \tilde{x}_t$. Thus, fluctuations of the variable x_t around its steady state are additively decomposed as

$$x_t - \bar{x} = (\bar{x}_t - \bar{x}) + (\tilde{x}_t - \bar{x}_t) + (x_t - \tilde{x}_t), \quad (1)$$

where the first bracket captures long-run fluctuations around the steady state, the second bracket captures medium-run fluctuations around the long-run trend, and the third bracket captures short-run fluctuations around the medium-run trend. Importantly, ε_t^x also captures the mismeasurement of variables present in the data, in addition to proxying short-run fluctuations.

3.2 Expectations

It is assumed that expectations are backward-looking ($E_{t-1} \tilde{\pi}_t^p$), and these are modeled with exponential smoothing.⁹ In particular, for a generic random variable x_t , the corresponding smoothed variable $\tilde{x}_t$ is defined as the weighted average of $\tilde{x}_{t-1}$ and x_{t-1} , with $\tilde{x}_t = (1 - \theta) \tilde{x}_{t-1} + \theta x_{t-1}$, where $\theta \in (0, 1)$.

Backward substitution shows that the exponentially smoothed variable is a weighted average of past realizations of the random variable x_t , with $\tilde{x}_t = \theta \sum_{j=1}^{\infty} (1 - \theta)^{j-1} x_{t-j}$. Additionally, expectations are adaptive, as past overshoot of x_t above expectations leads to an upward revision of expectations. In the case of price inflation in the medium run, the exponential smoothing implies the following relation:

$$E_{t-1} \tilde{\pi}_t^p = (1 - \theta) E_{t-2} \tilde{\pi}_{t-1}^p + \theta \tilde{\pi}_{t-1}^p. \quad (2)$$

⁹For a link between exponential smoothing and state-space representation, see Hyndman et al. (2008).

Thus, expectations are modeled as a function of past data and do not incorporate survey-based measures, as for example in Bernanke and Blanchard (2023).

3.3 Goods Market

The analysis adopts the standard New Keynesian logic, positing that medium-run inflation is determined by the output gap, the real wage gap, and expected inflation. Hence, the following empirical specification is adopted:

$$\begin{aligned}\tilde{\pi}_t^p - \bar{\pi}^p &= \beta^p (E_{t-1} \tilde{\pi}_t^p - \bar{\pi}^p) + a_{p\Delta y} \Delta \hat{y}_t \\ &+ a_{py} \hat{y}_{t-1} + a_{p\varpi} \hat{\varpi}_{t-1} + v_t^p.\end{aligned}\quad (3)$$

The equation is mostly standard but allows for a lag structure, as expanding the difference allows for the output gap to enter the equation both contemporaneously and with a lag.¹⁰ The difference serves a purpose: the output gap is a very low-frequency variable, and adding the difference brings it to a scale similar to inflation. Furthermore, the real wage gap is introduced with a lag as well. As the inflation Phillips curve effectively represents the supply curve for goods and services, v_t^p in this context can be interpreted as supply shocks that account for the residual variation of inflation. These shocks are assumed to follow a first-order autoregressive process.

Consistent with the rationale of the New Keynesian Phillips curve, higher demand results in the average markup falling short of its desired level, prompting firms to adjust prices upward. Similarly, a larger real wage gap leads to a below-desired average markup, causing firms to increase prices. In addition to the economic logic of the New Keynesian model, from a statistical standpoint, the lagged real wage gap acts as an error-correction term, exerting upward pressure on inflation when the gap is positive.

Next, the demand for goods and services is modeled as follows:

$$\hat{y}_t = \rho_y \hat{y}_{t-1} - a_{yp} (\tilde{\pi}_t^p - \bar{\pi}^p) + a_{ye} (E_{t-1} \tilde{\pi}_t^p - \bar{\pi}^p) + v_t^y. \quad (4)$$

¹⁰ As a theoretical reference, in the canonical New Keynesian model outlined by Galí (2015) in Chapter 7, the inflation Phillips curve adopts the following form: $\pi_t^p = \beta E_t \pi_{t+1}^p + \lambda_{py} \hat{y}_t + \lambda_{p\varpi} \hat{\varpi}_t$.

The equation assumes that the output gap is a negative function of current inflation and a positive function of expected inflation. This specification contrasts with Bernanke and Blanchard (2023), who do not model domestic slack endogenously. This distinction is crucial because, in the medium run, output is highly susceptible to changes in inflationary pressures, either directly through demand as modeled in this paper, or indirectly via the monetary policy stance. The demand equation above stems from rewriting the real interest rate gap, using the Fisher identity and the following interest rate rule:¹¹

$$i_t^p = r_t^n + \bar{\pi}^p + \frac{a_{yp}}{a_{ye}} (\tilde{\pi}_t^p - \bar{\pi}^p) + e_t^i.$$

With the assumed policy rule, the natural rate does not affect the output gap, conditional on the interest rate rule. A key advantage of this approach is that it avoids the need to model the nominal interest rate explicitly, which is empirically challenging given its persistent downward trend in the euro area over the sample period. This trend renders the nominal rate similar to a non-stationary series, complicating estimation. However, a potential drawback is that the model may understate the role of natural rate movements in the business cycle. Rules of other forms that do not neutralize the natural rate require the explicit modeling of the nominal interest rate.¹² In the present specification, residual variation in aggregate demand reflects both exogenous monetary policy shifts and other demand disturbances, which cannot be separately identified. Finally, the composite error v_t^y follows a first-order autoregressive process.

3.4 Labor Market

Moving on to wages, the growth rate of nominal wages is determined by expectations, as well as the employment gap. The use of employment instead of unemployment allows for accommodating features

¹¹As a theoretical reference, in the canonical New Keynesian model described by Galí (2015) in Chapter 7, the corresponding equation takes the form $\hat{y}_t = E_t \hat{y}_{t+1} - \lambda_{yr} (i_t - E_t \pi_{t+1}^p - \tilde{r}_t)$.

¹²Regarding interest rate smoothing, the nominal interest rate can be expressed as a function of medium-run inflation, noted for its stickiness. Given a rule of the form $i_t = \rho i_{t-1} + (1 - \rho)\pi_t$, where π_t represents current headline inflation, solving for i_t implies that interest rates depend on both current and past inflation rates. This is encapsulated by the concept of medium-run inflation.

such as trends that are present in the data.¹³ Thus the wage Phillips curve is specified as follows:

$$\begin{aligned}\tilde{\pi}_t^w - \bar{\pi}_t^w &= \beta^w (E_{t-1} \tilde{\pi}_t^p - \bar{\pi}^p) + a_{w\Delta n} \Delta \hat{n}_t + a_{wn} \hat{n}_{t-1} \\ &\quad - a_{w\varpi} \hat{\varpi}_{t-1} + v_t^w,\end{aligned}\tag{5}$$

where $\bar{\pi}_t^w = \Delta \bar{\varpi}_t + \bar{\pi}^p$ and $\hat{n}_t$ measures the employment gap.¹⁴ Similar to the logic presented for Equation (3), both the change in the employment gap and the lag of the employment gap are present in the equation. Additionally, the lag of the real wage gap is introduced to provide greater empirical flexibility.¹⁵ It is important to note that in this specification, it is the expectation of price inflation that influences wage setting, diverging from the approach of using expectations of wage growth as presented by Galí (2015). This approach aligns with the empirical methodology adopted by Bernanke and Blanchard (2023). As before, v_t^w explains the residual variation of the wage inflation, and follows a first-order autoregressive process.

The wage equation effectively captures the labor supply curve. The relation states that a decline in labor demand, for a given level of labor supply, reduces employment. The resulting excess supply of labor causes wages to fall. Similarly, an increase in labor supply for a given level of labor demand results in higher unemployment and lower wages. From a statistical perspective, the lagged real wage gap functions as an error-correction term, exerting upward pressure on wages when the gap is negative.

The next key equation links the employment gap to the output gap and the real wage gap, as well as its own lag:

$$\hat{n}_t = \rho_n \hat{n}_{t-1} + a_{ny} \hat{y}_t - a_{n\varpi} \hat{\varpi}_t + v_t^n.\tag{6}$$

¹³Using the definition of the unemployment rate combined with the labor force, the employment level can be expressed as $N_t = (1 - u_t)L_t$. This indicates a negative correlation between the unemployment rate and the employment gap.

¹⁴For a theoretical reference, in the canonical New Keynesian model described by Galí (2015) in Chapter 7, the counterpart equation takes the form $\pi_t^w = \beta E_t \pi_{t+1}^w - \lambda_{wu} \hat{u}_t$. The deviation from the canonical model involves using the employment gap instead of the unemployment rate gap, leading to a reversal of the sign on the relevant variables.

¹⁵For example, in the New Keynesian model described by Galí (2015) in Chapter 7, the real wage gap influences wages through labor demand only.

As the variables on the right-hand side are primarily slow-moving, in the current specification all factors on the right-hand side enter contemporaneously.¹⁶ The equation suggests that a higher demand for goods and services results in a greater demand for labor, which, given a constant level of labor supply, increases the employment gap. Conversely, a larger real wage gap is positively associated with a higher labor supply. For a constant level of labor demand, this reduces the employment gap. As before, v_t^n explains the residual variation of the employment gap, and is modeled as a first-order autoregressive process.

Equations (2)–(6) together with the equations for variables in the long run and steady-state constraints form the system.

4. Selection of Priors

The system comprises 30 parameters that need to be estimated. It also includes 17 unobservable endogenous variables. Each of the four observable variables has 114 observations. Except for the steady-state inflation, which is fixed at 0.5, all other parameters are estimated. Given the structure of the empirical model, Bayesian estimation methods complemented by Kalman filtering techniques are essential. The approach adopted is primarily data driven due to limited prior knowledge about the magnitudes of the parameters, except for their signs.

4.1 Shocks

To gain insights into the temporal dimension, $\ddot{x}_t$ is defined as the four-quarter moving average for $x_t \in (\pi_t^p, \pi_t^w, \Delta n_t, \Delta y_t)$, serving as the initial proxy for the medium-run trend.¹⁷ It is important to emphasize that these are used solely for the initial estimation of the standard deviations of the short-run and medium-run shocks. They are not incorporated into system estimation as a medium-run

¹⁶For a theoretical reference, in the canonical New Keynesian model described by Galí (2015) in Chapter 7, the counterpart equation takes the form $\hat{u}_t = -\lambda_{uy}\hat{y}_t + \lambda_{u\varpi}\hat{\varpi}_t$. The deviation from the canonical model involves using the employment gap instead of the unemployment rate gap, leading to a reversal of the sign on the relevant variables.

¹⁷Note that, at an annual frequency, this is equivalent to seven quarterly observations.

Table 2. Priors for Shocks

	(1) π_t^p	(2) π_t^w	(3) Δn_t	(4) Δy_t
$SD(\varepsilon_t^x)$	0.217 (0.018)***	0.277 (0.030)***	0.200 (0.029)***	0.562 (0.077)***
$SD(e_t^x)$	0.081 (0.007)***	0.105 (0.008)***	0.072 (0.010)***	0.211 (0.027)***
$\frac{SD(\varepsilon_t^x)}{SD(e_t^x)}$	2.669 (0.235)***	2.632 (0.264)***	2.778 (0.365)***	2.660 (0.290)***
Observations	110	110	110	110
Note: The table provides the results from a state-space estimation, as explained in Subsection 4.1. These are then used as priors in the Bayesian estimation. Robust standard errors are shown in parentheses. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels, respectively.				

trend. Next, within the framework of a state-space representation, an observation equation is defined as follows:

$$x_t = \ddot{x}_t + \varepsilon_t^x, \quad (7)$$

where ε_t^x is unobservable and captures short-run fluctuations, and $\ddot{x}_t$ serves as a proxy for the medium-run trend. To obtain an initial estimate of the standard deviation for medium-run shocks, it is further assumed that

$$\ddot{x}_t = \eta^x \ddot{x}_{t-1} + \tau_t^x, \quad (8)$$

where τ_t^x is unobservable and follows a first-order autoregressive process with a disturbance denoted by e_t^x .¹⁸ These two equations are estimated for each x_t . The results of the estimates of σ_{ε^x} and σ_{e^x} with their corresponding standard errors are shown in rows 1 and 2 of Table 2. A normal distribution with mean and standard deviation taken from the first row of the table serves as the priors for standard deviation of short-run shocks denoted by σ_{ε^x} .

¹⁸The two equations can be estimated individually without the state-space representation. However, with the state-space representation, it is relatively easy to estimate the standard deviation of ε_t^x and e_t^x .

The final row of the table presents the ratio $\sigma_{\varepsilon^x}/\sigma_{e^x}$ along with its associated standard error. Although there is no theoretical basis for this pattern, the ratio seems to remain relatively stable across the variables. Hence, the standard deviation of medium-run shocks are defined by $\sigma_{e^x} = \lambda\sigma_{\varepsilon^x}$, where $\lambda = \sum(\sigma_{\varepsilon^x}/\sigma_{e^x})/4$. The prior for the λ parameter is also set to a normal distribution with a mean of 2.7 (the arithmetic average of $\sigma_{\varepsilon^x}/\sigma_{e^x}$ across the four variables). Similarly, the prior value for the standard deviation of this distribution is derived from the arithmetic average of the standard errors across the four variables (0.3 in this case). This approach narrows the parameter space, facilitating the estimation of 30 parameters instead of 33.


4.2 Elasticities

Turning to sensitivities, they are more likely to fall within the interval of zero to one. Using uniform priors is strongly discouraged, so curved priors are adopted instead (Gelman and Yao 2021). Additionally, in modal estimation, it is crucial to eliminate the boundary problem. A gamma prior appears suitable for this purpose. To determine the optimal generic parameterization for the gamma distribution, 100 million draws from a truncated $N \sim (0.5, 0.5)$ distribution are initially generated. Parameters of the gamma distribution are then estimated based on this sample. The resulting estimates for the shape and scale parameters are 1.885 and 0.341, respectively. These parameters infer that the mean, standard deviation, and mode of the gamma distribution are 0.644, 0.469, and 0.576, respectively. These values are set as priors across all elasticities that are not bounded between zero and one. Panel A of Figure 2 displays the shape of this distribution.¹⁹

For parameters bounded between zero and one, a similar approach is undertaken. Drawing from 100 million samples of a truncated $N \sim (0.5, 0.5)$ distribution, the parameters of a beta distribution are then estimated to achieve the best representation. The resulting alpha and beta parameters for the beta distribution are both estimated at 1.173. This yields a mean of 0.5 and a standard

¹⁹While the gamma distribution is not constrained on the right, only the zero-to-one interval is plotted for clarity.

Figure 2. Prior Distributions



Note: The figure illustrates the prior distributions. While the gamma distribution is not constrained on the right, only the zero-to-one interval is plotted for clarity.

deviation of 0.273. These values are subsequently used as priors for this set of parameters. Panel B of Figure 2 displays the shape of this distribution.


4.3 Long-Run Growth

Finally, the long-run properties of the growth rate are extracted by using several low-frequency trigonometric weighted averages (Müller and Watson 2015). In particular, the cosine projections are constructed as the fitted values obtained by regressing Δy_t onto a constant term and $\sqrt{2} \cos(j(t - 0.5)/T)$ for $j = 1, \dots, 12$. Next, using the resulting series, a simple ARIMA(1,0,0) is estimated to determine the standard deviation of the long-run growth rate of output, the persistence parameter of the process, and the associated standard errors. The point estimate for the standard deviation of the long-run growth rate of output stands at 0.057, with an associated standard error of 0.013 (statistically significant at 1 percent). These estimates serve as priors for σ_χ with $N \sim (0.057, 0.013)$. The estimated persistence parameter is 0.988 with a corresponding standard error of 0.005. These estimates are used as priors for ρ_χ with $N \sim (0.988, 0.005)$.

Throughout the estimation, all normal priors are truncated on the left. The estimation begins by finding the mode of the posterior distribution and then implementing the Metropolis-Hastings algorithm. A single chain is generated, consisting of 50 million draws, with an acceptance rate of 33 percent. After discarding the first 25 million draws, the convergence diagnostics signal convergence of the chain. The estimation procedure is executed using Dynare (Adjemian et al. 2022).

5. Estimation Results

The Dynare identification toolbox is utilized to verify the identification of parameters in the model calibrated at the mode. The three built-in tests reveal that all parameters are identified in the Jacobian of the (i) steady-state and reduced-form solution matrices; (ii) mean and spectrum (Qu and Tkachenko 2012); and (iii) first two moments (Iskrev 2010). However, the data sample used for estimation is notably short, especially considering the

Figure 3. Posterior Distributions

Note: The figure illustrates the prior and posterior distributions. The prior distribution is presented by the solid black line, while the posterior distribution is presented by the dotted blue line.

number of parameters and unobservable variables that require estimation. Consequently, in certain instances, some parameters are weakly identified (see Figure 3). Further investigation is necessary to address these limitations. Nonetheless, the impulse response functions and the historical shock decomposition appear reasonable, suggesting that the estimated parameters are largely plausible.

5.1 *Parameters*

Before examining the parameters of the medium-run system, it is instructive to consider the implications of the estimated parameters in the steady state.²⁰ The estimates imply the following steady-state patterns on an annual basis. The rate of economic expansion, measured by the growth rate of GDP, stands at $\Delta\bar{y} = 1.34$ percent, while the growth rate of employment is approximately $\Delta\bar{n} = 0.86$ percent. From these estimates, per capita output growth rate in the steady state is deduced to be around $\Delta\bar{y} - \Delta\bar{n} = 0.47$ percent. Consequently, the growth rate of real wages is $\Delta\bar{w} = 0.47$ percent. Inflation, gauged by the logarithmic change in the GDP deflator, is fixed at $\bar{\pi}^p = 2$ percent in the steady state. Based on these estimates, the steady-state growth rate of nominal wages is deduced to be $\bar{\pi}^w = 2.47$ percent.

Table 3 displays the estimation results for the primary parameters of interest, grouped by equation. Interpreting the raw coefficients of dynamic models directly can be misleading (e.g., coefficients of a VAR model). This is because these models involve complex interrelationships between variables over time, and the coefficients simply capture a snapshot of those relationships at one point in time. In such models, the interpretation of coefficients as marginal effects (as in linear regression) does not hold. Instead, a more intuitive way to understand the dynamics of the model is through impulse response functions. Nevertheless, the text below provides an overview of estimated parameters, focusing on the values at mode and corresponding standard deviations.

The parameter θ quantifies the extent to which inflation expectations are revised. A value closer to one suggests a quicker adjustment of expectations. This parameter has an estimated value of roughly 0.47, accompanied by a standard deviation of 0.31. Such a coefficient highlights a significant role historical inflation plays in shaping expectations. For example, if past inflation exceeds prior expectations by a percentage point, inflationary expectations typically rise by approximately 47 basis points. The estimated parameter also

²⁰For conciseness, these estimates are not presented but are available upon request.

Table 3. Posterior Estimation of Parameters

	Distr.	Prior		Posterior				
		Mean	St. Dev.	Mode	Mean	St. Dev.	5 Pct.	95 Pct.
θ	Beta	0.500	0.273	0.465	0.434	0.309	0.000	0.786
β^p	Beta	0.500	0.273	0.403	0.419	0.358	0.023	0.751
$a_p\Delta y$	Gamma	0.644	0.469	0.147	0.167	0.076	0.041	0.285
a_{py}	Gamma	0.644	0.469	0.070	0.099	0.054	0.012	0.179
$a_{p\overline{w}}$	Beta	0.500	0.273	0.142	0.162	0.132	0.000	0.309
ρ_{v^p}	Beta	0.500	0.273	0.812	0.787	0.137	0.603	0.977
β^w	Beta	0.500	0.273	0.118	0.361	0.255	0.000	0.726
$a_w\Delta n$	Gamma	0.644	0.469	1.419	1.431	0.244	1.029	1.830
a_{wn}	Gamma	0.644	0.469	0.139	0.164	0.105	0.016	0.301
$a_{w\overline{w}}$	Beta	0.500	0.273	0.945	0.853	0.117	0.697	1.000
ρ_{v^w}	Beta	0.500	0.273	0.928	0.551	0.110	0.125	0.991
ρ_n	Beta	0.500	0.273	0.233	0.238	0.202	0.003	0.432
a_{ny}	Gamma	0.644	0.469	0.620	0.740	0.153	0.473	0.995
$a_{n\overline{w}}$	Gamma	0.644	0.469	0.383	0.535	0.257	0.122	0.920
ρ_{v^n}	Beta	0.500	0.273	0.472	0.514	0.586	0.112	0.962
ρ_y	Beta	0.500	0.273	0.793	0.746	0.128	0.557	0.953
a_{yp}	Gamma	0.644	0.469	0.162	0.371	0.170	0.014	0.725
a_{ye}	Gamma	0.644	0.469	0.203	0.389	0.208	0.016	0.759
ρ_{v^y}	Beta	0.500	0.273	0.796	0.762	0.119	0.598	0.937

Note: The posterior distribution is obtained using the Metropolis-Hastings algorithm. The estimations are performed in Dynare with a single chain consisting of 50 million draws, an acceptance rate of 33 percent, and a 50 percent burn-in.

suggests that contributions of past inflation to expectations become negligible after the first four quarters.

Regarding the inflation Phillips curve, coefficient β^p captures the responsiveness of inflation to its expectations. The estimated value is about 0.40, with a standard deviation of 0.36. The extent to which inflation correlates with the output gap is represented by the $a_{p\Delta y}$ and a_{py} parameters. The former shows how inflation reacts to immediate changes in the output gap, given a constant demand. With an estimate of 0.15 and a standard deviation of 0.08, this suggests that fluctuations in the output gap can notably influence inflation. On the other hand, a_{py} , which indicates how past output gaps affect current inflation, has an estimated value of 0.07 and a standard deviation of 0.05. This effect appears more muted, possibly due to inflation and the output gap having different amplitudes of movement.

Parameter $a_{p\varpi}$ shows the impact of the lagged real wage gap on inflation. In the current empirical framework, it can be used to interpret the correction of inflation toward its long-run value. For example, when the real wage gap is positive, the correction requires an increase in prices. The estimated value for this parameter is 0.14, coupled with a standard deviation of 0.13. Finally, the persistence of price shocks is estimated at 0.81 with a 0.14 standard deviation.

In the wage equation, coefficient β^w reflects wage sensitivity to expected inflation, estimated at 0.12 with a standard deviation of 0.26. The $a_{w\Delta n}$ parameter measures immediate wage response to changes in the employment gap, conditional of keeping the other factors fixed. Estimated at 1.42, with a standard deviation of 0.24, it suggests that employment gap variations can strongly affect wage growth. On the other hand, the a_{wn} parameter denoting the impact of past employment gaps on current wage growth has an estimate of 0.14 and a standard deviation of 0.10. The two estimates suggest that fluctuations of medium-run employment around its long-run trend are more tightly linked to movements in wages, compared to the relationship between inflation and the output gap.

Similar to inflation, the $a_{w\varpi}$ parameter measures error correction and indicates the required wage adjustment in light of deviations from long-run equilibrium. Its estimated value is 0.95 with a standard deviation of 0.12. The estimate is considerably higher than the one for inflation, and the standard deviation at the mode is smaller, suggesting that error correction significantly influences

euro-area wage dynamics. Finally, the estimated persistence for wage shocks is quite significant at 0.93, with a standard deviation of 0.11, primarily reflecting developments toward the end of the sample.

Turning to labor demand, the employment gap demonstrates moderate persistence, with an estimated parameter of 0.23 and a standard deviation of 0.20. The sensitivity of the employment gap to the output gap, represented by the coefficient a_{ny} , is estimated at 0.62 with a standard deviation of 0.15. Meanwhile, the sensitivity to the real wage gap, captured by the coefficient $a_{n\varpi}$, is estimated at 0.38, accompanied by a standard deviation of 0.26. Additionally, the persistence of labor demand shocks is estimated at 0.47 with a standard deviation of 0.59.


Lastly, aggregate demand demonstrates a high level of persistence, with a value of 0.80 and a standard deviation of 0.13. The persistence in demand is likely indicative of habit formation among consumers. The sensitivity of demand to current inflation, represented by a_{yp} , is estimated at 0.16 with a standard deviation of 0.17. Sensitivity to expected inflation, indicated by a_{ye} , stands at 0.20 with a standard deviation of 0.21. Demand shocks are also notably persistent, having a coefficient of 0.80 and a standard deviation of 0.12.

Table 4 summarizes the results for the estimated standard deviations of short-run and medium-run shocks. The estimated standard deviations for short-run fluctuations in inflation, wages, employment, and output are, respectively, 0.21 (0.01), 0.22 (0.02), 0.11 (0.01), and 0.52 (0.04)—with standard errors in parentheses. The implied standard deviations for medium-run shocks for these variables can be determined by dividing the short-run standard deviation by the parameter λ . In the same order, these are 0.11, 0.12, 0.06, and 0.29.

5.2 *Signal and Noise*

The blue line in Figure 4 represents the actual quarterly series, while the red line depicts the extracted medium-run trend for the four observed variables. Panel A presents the results for inflation. The derived medium-run trend is noticeably smoother, effectively reducing pronounced volatility. Panel B displays the results for wages. The medium-run trend for wages exhibits a slightly higher degree of

Figure 4. Signal and Noise



Note: The figure displays the actual series along with its corresponding medium-run trend. The last in-sample observation is for 2023:Q3.

Table 5. Variance Decomposition

	e_t^p	e_t^w	e_t^y	e_t^n	e_t^x
$\tilde{\pi}_t^p$	54.5	3.4	42.1	0.1	0.0
$\tilde{\pi}_t^w$	43.6	6.9	44.0	4.6	1.0
$E_{t-1}\tilde{\pi}_t^p$	52.5	3.7	43.8	0.0	0.0
$\hat{y}_t$	0.1	0.0	99.9	0.0	0.0
$\hat{n}_t$	0.7	1.0	98.1	0.2	0.0
$\widehat{\omega}_t$	23.0	40.2	34.8	2.0	0.0
Note: The table presents the variance decomposition of key macroeconomic variables, with parameters calibrated to the estimated mode.					

volatility compared to inflation. This difference primarily arises from the modeling assumptions, where long-run nominal wage growth is determined by a mix of equilibrium long-run real wages and a steady-state level of inflation.

Both inflation and the growth rate of nominal wages display consistent patterns, with values nearing the steady-state levels before the Global Financial Crisis. After the crisis, both measures declined below their steady-state values, recovering gradually up to the onset of the COVID-19 shock. Following this event, both indicators register values that significantly surpass the steady state.

A somewhat similar pattern emerges for the growth rates of output and employment. The medium-run trend effectively filters out short-run fluctuations. Both series witnessed declines following the Global Financial Crisis and the COVID-19 shocks, with subsequent recoveries. The medium-run trends for both output and employment growth then display renewed momentum.

6. Decompositions

6.1 Variance Decomposition

Table 5 provides the variance decomposition that sheds light on the contribution of different shocks to fluctuations in key economic variables. For inflation, the price shock is the predominant contributor, accounting for 54.5 percent of the variations. Following closely,

the demand shock explains 42.1 percent, underscoring that exogenous shifts in aggregate demand significantly influence inflationary dynamics. Other shocks make a minor impact. Unsurprisingly, expected inflation is explained by the same shocks with somewhat similar contributions. Wage growth is also influenced significantly by the price shock at 43.6 percent, while the demand shock accounts for 44.0 percent. The wage shock and labor demand shock play smaller roles, contributing 6.9 and 4.6 percent, respectively.

When examining the output gap and employment gap, both variables are almost entirely shaped by exogenous shifts in aggregate demand, with large contributions of 99.9 and 98.1 percent, respectively. The influences of other shocks on these variables are minimal. The real wage gap sees a wider influence from the exogenous factors. In particular, the wage shock accounts for 40.2 percent in wage variation, while the demand shock and price shock contribute around 34.8 and 23.0 percent, respectively. This underscores the pivotal roles these shocks play in determining fluctuations in real wages.


Overall, the demand shock significantly influences the key economic variables. Conversely, the price shock emerges as the primary force behind inflation, wage growth, and the real wage gap, while wage shocks explain more recent variation in the wage gap. Other shocks provide supplementary insights.

6.2 Shock Decomposition

Figure 5 presents the results of the historical shock decomposition over the sample period. The subsequent text focuses on the more recent period from 2019:Q1 to 2023:Q3.

Panel A presents the decomposition for medium-run inflation. From 2019:Q1 to 2019:Q3, inflation hovered just below the ECB target, and was on target on 2019:Q4. This period was marked by positive demand shocks counteracted by negative price shocks. Between 2020:Q1 and 2020:Q3, negative demand shocks pushed medium-run inflation well below its steady state. Starting in 2020:Q4, price shocks took hold and, combined with positive demand shocks, propelled medium-run inflation above the steady-state by 2021:Q2. Subsequent to this, both price and demand shocks have been driving inflation upward.

Figure 5. Shock Decomposition



Note: The figure illustrates the historical shock decomposition of key economic variables. The last in-sample observation is for 2023:Q3. The variables are measured relative to the steady state. All gaps have a steady-state value of zero. The quarterly steady-state value for inflation relative to the previous quarter is $\bar{\pi}^p = 0.5$. For wage growth, the quarterly steady-state value relative to the previous quarter is $\bar{\pi}^w = 0.62$. Finally, the quarterly steady-state value for output growth compared to the previous quarter is $\Delta\bar{y} = 0.33$, while the corresponding value for employment growth is $\bar{n} = 0.22$.

Panel B details the decomposition for the medium-run growth rate of nominal wages. Starting in 2019:Q1, there was a declining trend in the wage growth, primarily attributed to price shocks. In 2020:Q1, a sharp decline in wage growth occurred, largely influenced by demand shocks. From 2021:Q2, a mix of positive price and demand shocks drove the growth rate of nominal wages notably above its steady state.

Panels C and D outline the decomposition for the output gap and the employment gap, respectively. These gaps display highly similar patterns, mainly because the employment gap is substantially explained by the output gap. Both gaps are majorly influenced by shocks to aggregate demand. Between 2019:Q1 and 2019:Q4, demand shocks kept the gaps positive. However, the stretch from 2020:Q1 to 2021:Q1 was dominated by negative shocks, keeping both gaps negative. From 2021:Q3 onward, positive demand shocks have elevated both gaps above their pre-pandemic levels. Interestingly, while price shocks have slightly mitigated the output gap, they have contributed positively to the employment gap. This latter effect has been mediated by the real wage gap.

Panel E shows the decomposition for the real wage gap. Prior to the pandemic, the gap was predominantly positive, spurred by positive price and demand shocks. From 2019:Q4 to 2020:Q3, negative demand shocks shifted the gap to negative territory. Between 2021:Q1 and 2022:Q1, positive output shocks moved the gap back into positive territory. However, even though demand shocks remained positive thereafter, dynamics of the wage gap were predominantly influenced by significant price shocks, plunging the real wage gap into negative territory.


Finally, although the cyclical variation in GDP growth rate is primarily driven by medium-run shocks, the growth rate has consistently remained below the steady state since the Global Financial Crisis. This is mainly due to the below-steady-state growth rate of the long-run output (panel F).

While this analysis is informative, it does not provide deep insights into factor contributions. For example, the price shock affects inflation, inflation expectations, and the wage gap. The current decomposition does not explain how each variable is associated with the dynamics of the primary variables under consideration. Consequently, a more detailed factor decomposition is presented below.

6.3 Factor Decomposition

Figure 6 presents the results of the factor decomposition over the sample period. While the factor decomposition is informative, it is

Figure 6. Factor Decomposition



Note: The figure illustrates the factor decomposition of key economic variables. The last in-sample observation is for 2023:Q3. The variables are measured relative to the steady state. All gaps have a steady-state value of zero. The quarterly steady-state value for inflation relative to the previous quarter is $\bar{\pi}^p = 0.5$. For wage growth, the quarterly steady-state value relative to the previous quarter is $\bar{\pi}^w = 0.62$. Finally, the quarterly steady-state value for output growth compared to the previous quarter is $\Delta \bar{y} = 0.33$, while the corresponding value for employment growth is $\bar{n} = 0.22$.

important to recognize that all factors are interrelated. The subsequent text focuses on the more recent period from 2019:Q1 to 2023:Q3.

Panel A shows the results of the decomposition for medium-run inflation. The period spanning from 2019:Q1 to 2019:Q4 was largely characterized by a positive output gap, positive real wage

gap, and the recovery of inflation expectations toward their steady-state value. These factors were somewhat counterbalanced by negative price shocks. As a result, inflation hovered slightly below its steady-state value, with annualized medium-run inflation reaching close to 2 percent in 2019:Q4. The outbreak of the pandemic in 2020 brought about significant changes. In 2020:Q1 and 2020:Q2, a sharp contraction in aggregate demand and falling inflation expectations led inflation to dive below its steady-state value. From 2020:Q4, notable price shocks arose. Along with rising inflation expectations and a positive output gap, inflation has moved well above its steady state.

Panel B shows the results of the decomposition for medium-run growth rate of nominal wages. During 2019, the labor market was tight, characterized by a positive employment gap. This had a favorable effect on wage growth, although the effect gradually diminished. When combined with the correction exerted on wages by the positive real wage gap, the trajectory of nominal wages trended downward from its steady-state value. The pandemic then induced a sharp employment drop, which in turn suppressed wage growth. From 2021:Q2 onward, labor markets displayed signs of tightness and inflation expectations began to climb, both of which contributed to increasing wage growth. This upward trend was partially counteracted by the positive real wage gap that persisted from 2021:Q2 to 2022:Q1. From 2022:Q2 onward, the upward corrective pressure from the negative real wage gap, labor market tightness, and mounting inflation expectations have driven wage growth significantly above its steady-state level.

Panel C displays the decomposition of the output gap. In 2019, the economy was overheating, evidenced by a positive output gap of approximately 1.5 percent. The periods 2020:Q1 and 2021:Q2 experienced a significant shock due to the COVID-19 pandemic. This led to a decline in the output gap, pushing it into negative territory, where it lingered from 2020:Q2 through to 2021:Q1. As the economy began to recover from the pandemic, a series of positive demand shocks propelled the output gap back into positive territory from 2021:Q3 onward. Coupled with the significant own persistence, the output gap has now surpassed its pre-pandemic level, reaching a value of 2.3 percent in 2023:Q3.

Panel D displays the results of the decomposition for the employment gap. The dynamics of the labor market, characterized by the employment gap, are intrinsically linked with those of aggregate demand, as represented by the output gap. From 2019:Q1 to 2019:Q4, a positive output gap applied upward pressure on labor demand. Yet, when the output gap contracted in 2020:Q1, it resulted in diminished labor demand. This negative employment gap persisted throughout the 2020:Q2–2021:Q2 period. Starting in 2021:Q3, labor demand began its recovery, and by 2023:Q3, the labor market had become tighter than before the onset of the pandemic.


Panel E illustrates the decomposition of the real wage gap. By its design, inflation exerts downward pressure on the real wage gap, whereas growth in nominal wages tends to elevate it relative to the long-run trend. The year 2019 was characterized by a positive gap, signifying that real wages exceeded their long-run trend. Nevertheless, there was an overarching downward trend during this time. Between 2020:Q1 and 2020:Q3, the real wage gap declined into negative territory. It then remained close to its long-run equilibrium from 2020:Q4 to 2021:Q1 before climbing back into positive territory, peaking in 2021:Q3. Since then, the trend of the real wage gap has been on a downward trajectory, turning negative in 2022:Q2 and declining further. This negative wage gap is likely to necessitate an adjustment in both nominal wages and prices.

Panel F displays the results for the medium-run growth rate of the output. Since this decomposition stems from the output gap identity, the growth rate essentially represents the slope of the output gap in panel C, conditional on the growth rate of long-run output. As mentioned earlier, beyond cyclical fluctuations, the medium-run growth rate has been below the steady-state level, primarily due to the subpar long-term growth rate of the output.

7. Dynamic Analysis

The variance decomposition suggests that most of variation comes from price and demand shocks. Consequently, impulse response functions for only these two shocks are discussed in the paper.²¹

²¹The remaining responses are available from the author upon request.

Figure 7. Response to a Shock e_t^P of Magnitude $\bar{\pi}^P$ 

Note: The figure illustrates the impulse response function of key economic variables to a shock e_t^P of magnitude $\bar{\pi}^P = 0.50$. The bars display the factor decomposition for each response variable. The dynamics are presented relative to the steady state. All gaps have a steady-state value of zero. The quarterly steady-state value for inflation relative to the previous quarter is $\bar{\pi}^P = 0.5$. For wage growth, the quarterly steady-state value relative to the previous quarter is $\bar{\pi}^w = 0.62$. Finally, the quarterly steady-state value for output growth compared to the previous quarter is $\Delta \bar{y} = 0.33$, while the corresponding value for employment growth is $\bar{n} = 0.22$.

7.1 A Price Shock

Figure 7 illustrates the response of the main variables after a shock to e_t^P with a magnitude of $\bar{\pi}^P = 0.5$. As the shock enters the inflation equation, it instantly elevates inflation by approximately 0.5 percentage points, resulting in a nearly identical negative real wage gap.

This negative real wage gap, in turn, boosts employment through the labor demand equation, as it enters the equation contemporaneously. Increased labor demand puts upward pressure on wages. At the same time, aggregate demand contracts due to the heightened inflation, which also enters the demand equation contemporaneously. These patterns are clearly observed in the first bar of each panel in the figure.


Following the initial shock, the variables begin a gradual adjustment back to their steady-state values. In panel A, the persistence of the shock maintains inflation above its steady-state value for an extended duration. Even as the influence of the shock wanes, inflation expectations rise, sustaining inflationary momentum. Meanwhile, the negative real wage gap exerts a corrective pressure on these inflationary forces in the opposite direction. Contributions from the output gap are negligible.

Panel B displays the breakdown of the wage response. Following the initial shock, the corrective force stemming from the real wage gap accelerates the growth rate of nominal wages. An increase in inflation expectations also contributes to wage growth.

Panel C offers the breakdown for the output gap. After its initial dip, the persistent nature of the gap depresses aggregate demand. Elevated inflation tends to dampen demand, while inflation expectations have the opposite effect. After the initial shock, pressure on aggregate demand arising from contemporaneous inflation dominates pressure from expected inflation. Over time, however, these factors largely neutralize each other.

Panel D depicts the employment response. The dynamics here are chiefly governed by the real wage gap. Since a negative real wage gap corresponds to diminished labor supply for a given labor demand, it raises the employment gap. Meanwhile, a decline in aggregate demand, and the associated reduction in labor demand, works to narrow the employment gap. Nevertheless, this latter effect is insufficient to counterbalance the forces arising from the real wage gap.

Panel E presents findings related to the real wage gap. This decomposition hinges on the real wage gap identity. After its initial decline, the decomposition reveals that the initial dominance of inflation recedes after the first few quarters. The subsequent dynamics are driven by the persistent narrowing of the gap towards its steady

Figure 8. Response to a Shock e_t^y of Magnitude $\Delta\bar{y}$ 

Note: The figure illustrates the impulse response function of key economic variables to a shock e_t^y of magnitude $\bar{y} = 0.33$. The bars display the factor decomposition for each response variable. The dynamics are presented relative to the steady state. All gaps have a steady-state value of zero. The quarterly steady-state value for inflation relative to the previous quarter is $\bar{\pi}^P = 0.5$. For wage growth, the quarterly steady-state value relative to the previous quarter is $\bar{\pi}^w = 0.62$. Finally, the quarterly steady-state value for output growth compared to the previous quarter is $\Delta\bar{y} = 0.33$, while the corresponding value for employment growth is $\bar{n} = 0.22$.

state. Lastly, the dynamics of the medium-run growth rate are solely influenced by the cyclical behavior of the output gap (panel F).

7.2 A Demand Shock

Figure 8 shows the response of the main variables resulting from a shock to e_t^y with a magnitude of $\Delta\bar{y} = 0.33$. As the shock enters

the demand equation, the output gap immediately rises by about 0.3 percentage points. This causes an instantaneous increase in both inflation and labor demand since the output gap affects inflation and labor demand contemporaneously. The rise in labor demand subsequently drives up wages. Since the initial wage increase surpasses that of prices, the real wage gap becomes positive. These patterns are distinctly noticeable in the first bar of each panel in the figure.

After this initial shock, the variables start a slow return to their steady-state values. Panel A provides a factor breakdown of the output gap. The early demand response is dominated by the positive shock but is soon surpassed by intrinsic dynamics of the demand itself. The persistent nature of the output gap continues to sustain its level considerably above the steady state for an extended period. Contributions from other factors are negligible.

Panel B depicts the employment response. In contrast to the price shock, the dynamics here are mainly determined by the evolution of the output gap. Higher demand for goods and services elevates labor demand, though to a lesser extent than the increase in the output gap itself. The response also reveals a slight persistence in labor demand influencing the employment gap. The corrective action from the real wage gap is negligible.

Panel C reveals the breakdown of the impulse response for inflation. The positive output gap remains the prevailing factor guiding the response. As inflation expectations pick up, they offset the dampening arising from convergence of the output gap to its steady state. Additionally, early in the response, the positive real wage gap exerts upward pressure on inflation. This turns negative toward the end of the response period, as the gap itself turns negative.

Panel D displays the breakdown of the wage response. Following the initial shock, the dynamics are dominated by tight labor markets, as indicated by the combined influence of current and past employment gaps. In the initial stages of the response, the corrective pressure stemming from the real wage gap tends to suppress wage growth. Toward the end of the horizon, as the wage gap turns negative, the correction stemming from the negative wage gap tends to sustain wage growth, while the tightness in the labor market effectively disappears. Finally, inflation expectations play a negligible role in wage dynamics following a demand shock.


Panel E presents findings related to the real wage gap. Like before, this breakdown relies on the real wage gap identity. After its initial surge, the factor decomposition shows that the initial dominance of wage growth diminishes after the first two quarters. The subsequent dynamics highlight a gradual reduction of the gap towards its steady state. Lastly, the cyclical behavior of the output gap solely determines the dynamics of the medium-run growth rate (panel F).

8. Forecasting Out of Sample

8.1 *Forecast Decomposition*

Figure 9 presents forecasts of key variables derived from the model calibrated with parameters set at the mode. The last in-sample observation is for 2023:Q3. Besides the forecast trajectory, the figure offers a decomposition of the forecasted variables by contributing factors. This exercise does not integrate any expert opinions. Forecasts rely on estimated parameters, current state estimates, and implicitly assume the absence of subsequent shocks. Thus, due to the stationarity of the system, forecasted outcomes tend to converge toward the steady state, irrespective of potential exogenous shocks that could arise, whether they are favorable or unfavorable.

Panel A portrays the forecast decomposition for inflation from 2023:Q4 to 2027:Q3. Variables shown are relative to the steady state, so a steady-state value of 0.5 must be added to inflation to obtain the point forecast estimate. The forecast for inflation is relative to the previous quarter. The point estimate for 2023:Q4 is 1.0 percent. Combining this with medium-run values for inflation at the start of 2023, and converting it into year-on-year average inflation, results in the annual inflation forecast for 2023 of 5.2 percent. Mean predictions for 2024, 2025, and 2026 are 3.8, 2.9, and 2.5 percent, respectively. Inflation expectations, followed by a positive output gap, are the primary drivers of this forecast. Although the corrective force of the real wage gap dampens inflation, its impact is small. Meanwhile, the projections for each quarter relative to the same quarter in the previous year are 4.8, 3.4, 2.7, and 2.3 percent for the years 2023 to 2026, respectively.

Figure 9. Out-of-Sample Forecasts

Note: The figure illustrates the out-of-sample forecast for key economic variables. The last in-sample observation is for 2023:Q3. The bars display the factor decomposition for each forecast variable. The dynamics are presented relative to the steady state. All gaps have a steady-state value of zero. The quarterly steady-state value for inflation relative to the previous quarter is $\tilde{\pi}^p = 0.5$. For wage growth, the quarterly steady-state value relative to the previous quarter is $\tilde{\pi}^w = 0.62$. Finally, the quarterly steady-state value for output growth compared to the previous quarter is $\Delta\tilde{y} = 0.33$, while the corresponding value for employment growth is $\tilde{n} = 0.22$.

Panel B provides the forecast decomposition for wage growth between 2023:Q4 and 2027:Q3. Given that the depicted variables are relative to the steady state, a value of around 0.6 should be added to obtain the point forecast estimate. The wage growth forecast is also relative to the previous quarter. The point estimate for 2023:Q4 is 1.1 percent. Combining this with medium-run wage growth values from the beginning of 2023, the year-on-year average forecast


for wage growth in 2023 is 5.3 percent. For 2024, 2025, and 2026, forecasts are 4.3, 3.5, and 3.0 percent, respectively. The forecast dynamics mainly stem from the corrective force of the real wage gap, with labor market tightness also playing a key role in the initial few quarters. Inflation expectations rank third. Meanwhile, the projections for each quarter, when compared to the same quarter of the previous year, are 5.1, 3.9, 3.3, and 2.9 percent for the years 2023 through 2026, respectively.

Panel C depicts the forecast decomposition for the output gap from 2023:Q4 to 2027:Q3. The output gap at the beginning of the forecast period is approximately 2.2 percent. Persistent aggregate demand chiefly drives the forecast dynamics. Current and expected inflation essentially balance each other in shaping demand dynamics, with a modest contribution from exogenous demand shocks. Panel D presents the forecast decomposition for the employment gap, which starts at 2.0 percent at the beginning of the period, underlining labor market tightness. The tightness is mainly attributed to the positive output gap. In addition, the negative real wage gap exerts minor upward pressure on the employment gap. Panel E displays the forecast for the real wage gap. Starting at approximately negative 0.3 percent at the beginning of the forecast horizon, it progressively returns to its steady-state value over the subsequent years.

Panel F illustrates the forecast for the quarter-on-quarter medium-run growth rate of output. At the start of the forecast horizon, a decline of around 13 basis points is projected relative to the steady-state growth rate of 0.33 percent. The trajectory is indicative of a substantial contraction and a subsequent gradual recovery. Finally, while the cyclical component of the growth rate shows signs of recovery, the medium-run growth rate is still constrained by the subpar long-run growth rate of output.

Focusing on the medium run provides a clearer understanding of the underlying trend in the data. However, this approach comes with drawbacks. By overlooking short-run fluctuations, important high-frequency events or shocks might be missed. Moreover, forecasts based solely on the medium run can be difficult to align with actual data. This might result in perceived inaccuracies in the forecast, even if the medium-run trend is captured relatively well. Nevertheless, Figure 10 provides a glance at the out-of-sample performance of

Figure 10. Out-of-Sample Forecast Performance



Note: The figure illustrates the out-of-sample forecast for key economic variables alongside their actual values, with forecasts performed as of 2023:Q3. $\tilde{\pi}_t^p$ represents the quarter-on-quarter forecast for medium-run inflation, while π_t^p presents the actual value of headline inflation. $\tilde{\pi}_t^w$ represents the out-of-sample quarter-on-quarter forecast for medium-run wage growth, while π_t^w presents the actual value of wage growth. $\Delta \tilde{y}_t$ represents the out-of-sample quarter-on-quarter forecast for the medium-run output growth rate, while Δy_t presents the actual value of the output growth rate. Finally, $\Delta \tilde{n}_t$ represents the out-of-sample quarter-on-quarter forecast for the medium-run employment growth rate, while Δn_t presents the actual value of the employment growth rate.

the forecasts. While the estimations are based on sample data from 1995:Q2 to 2023:Q3, data for 2023:Q4–2024:Q2 are also available.²²

Panels A for inflation and B for wage growth show that the three-quarter out-of-sample performance of the forecast is quite satisfactory, even though the comparison is between the medium-run trend and the reported headline values. For the growth rate of output, the discrepancy is larger, as can be seen in panel C. In particular, the projected medium-run growth rate of output is higher than the reported value of output growth rate. Interestingly, the reported value seems to be increasing while the medium-run value of the output growth rate shows a bit of decline, stabilizing around a somewhat similar value as the realized value of the output growth rate. Finally, panel D shows that the medium-run growth rate of employment is relatively stable, though some decline is visible. A somewhat similar pattern can be observed for the realized growth rate of employment, yet the discrepancy between the medium-run trend and the realized value is somewhat large. This, however, is not surprising, as Figure 6 shows that the medium-run trend is much smoother than the corresponding headline number, with clear patterns of misalignment in certain cases.

8.2 Discussion

8.2.1 On Projections

The ECB and Eurosystem staff continuously produce macroeconomic forecasts. Panel A of Figure 11 provides these projections for inflation together with the forecasts generated by this framework. ECB December 2023 projections reveal that inflation in the euro area, measured using the GDP deflator, is anticipated to average 5.6 percent in 2023.²³ This rate is expected to decline to 2.9 percent in 2024, further receding to 2.5 percent by 2025, and is at 1.9 percent in 2026. Panel B displays projections for wage growth.

²²At the time of writing, data for the deflator and compensation per employee are available until 2024:Q1. However, it is possible to construct synthetic series and expand them for an additional quarter using other highly correlated variables for which the data are available.

²³The ECB December 2023 projections were selected to ensure maximum comparability with the data used in this analysis.

Figure 11. Comparison



Note: In panel A, ECB projections for inflation are presented alongside forecasts produced by the current framework. Panel B provides ECB projections for wage growth, again paired with forecasts from the current framework. Panel C illustrates annualized inflation expectations at a quarterly frequency. Panel D offers annualized projections of model-generated inflation expectations, complemented by the annual inflation forecast. Panel E displays the annual output gap series generated by the model, alongside the annual output gap series from the IMF World Economic Outlook data set. Panel F shows the negative of the averaged quarterly model-generated employment gap over four quarters for a given year and the demeaned annual unemployment rate, sourced from the World Development Indicators data set.

For 2023, 2024, 2025, and 2026 the growth rate of nominal wages is expected to be 5.3, 4.6, 3.8, and 3.3. In contrast to these projections, the above forecast paints a somewhat grimmer picture for the medium-run inflation path.

As previously discussed, the projection for inflation from the model is primarily influenced by inflation expectations. Panel C displays annualized expectations at a quarterly frequency for model-generated inflation expectations. Given that model-generated expectations are based largely on the previous four lags of the underlying inflation, it is not surprising to see that these series are smooth and reflect the underlying trend. Running a simple ARIMA(1,0,0) on the series yields a persistence coefficient of 0.97 for the model-generated expectations. These estimates underscore the previously mentioned stickiness of expectations. Consequently, the sizable and enduring shock that occurred suggests a gradual decrease in expectations, leading to prolonged high inflation, as shown in panel D.

8.2.2 *On Shortcomings*

In general, differences in forecast outcomes often arise from variations in the methodologies used. Regarding the shortcomings of the current approach, the framework does not directly account for the prices of global intermediaries. Instead, it attributes these to the residual variation of the Phillips curve. Improved forecast accuracy might come from modeling this residual variation by analyzing exogenous factors directly. Adopting statistically robust values for these external processes, especially ones that resonate with expert assessments, could reinforce forecasting accuracy.

Another notable omission is the neglect of fiscal policy within the current empirical framework. It is either relegated to the residual variation of the output gap or implicitly incorporated into the gap itself. Given the significant impact of fiscal policy on aggregate demand, its modeling could provide a more nuanced understanding of inflationary trends and their underlying causes. For instance, while fiscal stimulus has been substantial in both the euro area and the U.S., it has arguably been stronger in the latter. This suggests that, compared to studies such as Bernanke and Blanchard (2023), the role of fiscal policy as a driver of inflation may be somewhat smaller in the euro-area context. Nonetheless, integrating fiscal policy could help improve forecasts. Prior work suggests that fiscal shocks can have non-negligible effects on inflation. For example, Cevik and Miryugin (2023) find that a one-standard-deviation decline in the budget balance-to-GDP ratio is associated with a

cumulative increase in headline inflation of around 0.7 percentage points over the medium run. Additionally, due to the significant number of parameters to be estimated, the external sector has been left out. The latter, however, can be tightly linked with medium-run inflationary developments (Galstyan 2019).

In addition, the recent wave of interest rate hikes has not yet been fully absorbed by aggregate demand. This suggests that the full implications of the ECB policy actions may not be entirely captured in the present data, which could skew estimates. It is worth noting that the recent surge in inflation is unparalleled in the euro-area history. Moreover, even with the aforementioned model enhancements, there is an underlying assumption that there have been no significant structural shifts in parameters. To account for potential shifts, continuous model reestimation is essential, possibly employing a rolling window to prioritize recent data. However, given the nature of macroeconomic data, this approach would result in a reduced data sample, necessitating stronger priors.

The insights derived from the dynamic responses and factor decomposition offer coherent results. Model-generated annual output gap series are somewhat similar to the output gap estimated by the International Monetary Fund, as shown in panel E of Figure 11. The employment gap series resemble the demeaned unemployment rate, as shown in panel F of Figure 11. Combined, these lend credence to the plausibility of the results. Nevertheless, it is essential to recognize that there is always room for potential pitfalls, whether arising from the choice of Bayesian priors or specific nuances in the specification of the model itself.²⁴

9. Conclusions

Understanding the dynamics of economic systems requires a multifaceted approach, encompassing the key variables across various time frames. By leveraging the logic of the New Keynesian model, this paper provides an empirical framework and analysis of the interrelations among inflation, wages, employment, and output, with a focus on the euro area.

²⁴For a discussion on the estimation of dynamic stochastic general equilibrium (DSGE) models, see Blanchard (2016).

After estimating the model using Bayesian techniques coupled with data-informed priors, the endogenous variables are decomposed into contributions from relevant factors, and dynamic analysis is conducted. The primary exogenous shocks that explain historical variance are price shocks and shocks to aggregate demand. The results highlight the significance of the wage gap in driving wage dynamics following a price shock, while the output gap appears to be more important in response to demand shocks. In the current inflationary environment, of particular interest are the real wage gap and inflation expectations.

Adjustments in the real wage gap offer a valuable perspective for understanding wage dynamics in the euro area. Statistically, the real wage gap functions as a corrective mechanism, ensuring alignment of wages and prices with the broader economic landscape. The forecasting exercise suggests the persistent significance of the real wage gap in the coming years. By 2023, given a declining real wage gap, there are inherent implications for necessary adjustments in both wages and prices. The corrective force of the gap will likely continue to influence nominal wages in the presence of tight labor markets.

Regarding inflation expectations, the estimates underscore their stickiness. Against this backdrop, the significant and persistent shock initially pointed to a gradual decline in expectations, raising concerns about a prolonged period of elevated inflation. However, the monetary policy response has been notably sharp, and inflation now appears to be gradually returning to target. This underscores the importance of timely and decisive monetary interventions in managing inflation expectations.

Further extensions of the framework are warranted in at least two dimensions. Firstly, recognizing that fiscal policy plays a pivotal role in shifting aggregate demand, its analysis would provide a clearer understanding of inflation trends and their determinants. Secondly, the current framework does not directly incorporate global intermediary prices into the system, relegating them to the residual variation of the Phillips curve. Accounting for these two elements could enhance the accuracy of inflation forecasts.

Appendix. Comparison to Bernanke and Blanchard (2023)

The 2023 Bernanke-Blanchard paper has drawn significant attention in policymaking circles. Given its prominence, a comparison of the approach in this paper with theirs is merited. Their model incorporates three primary equations. The first key equation is the wage equation and is given by²⁵

$$w_t - w_{t-1} = Ep_t - p_{t-1} - \alpha(Ep_{t-1} - p_{t-1}) + \beta(x_t - \alpha x_{t-1}) + z_t^\varpi. \quad (\text{A.1})$$

Using the notation of the current paper, this equation can be rewritten as

$$\pi_t^w = E\pi_t^p - \alpha(E\pi_{t-1}^p - \pi_{t-1}^p) + \beta\Delta x_t + \beta(1 - \alpha)x_{t-1} + z_t^\varpi. \quad (\text{A.2})$$

The second key equation is the price equation, which, in differences, is expressed as

$$p_t - p_{t-1} = (w_t - w_{t-1}) + (z_t^p - z_{t-1}^p), \quad (\text{A.3})$$

where z_t^p is captured by commodity prices, proxies for supply chains, and trends in productivity. The equation can be rewritten using the notation of the current paper as follows:

$$\pi_t^p = \pi_t^w + \Delta z_t^p. \quad (\text{A.4})$$

Finally, there are two types of inflation expectations. The short-run inflation expectations are given by

$$Ep_t - p_{t-1} = \delta\pi_t^* + (1 - \delta)(p_{t-1} - p_{t-2}), \quad (\text{A.5})$$

while the long-run expectations are expressed through

$$\pi_t^* = \gamma\pi_{t-1}^* + (1 - \gamma)(p_{t-1} - p_{t-2}). \quad (\text{A.6})$$

²⁵The equation is derived by assuming that the level of wages is determined according to $w_t = Ep_t + \varpi_t^A + \beta x_t$, where ϖ represents the real aspirational wage, intended to capture a “catch-up” of wages to its desired level, and x_t denotes the slack in the labor market. The real aspirational wage is, in turn, given by $\varpi_t^A = \alpha\varpi_{t-1}^A + (1 - \alpha)(w_{t-1} - p_{t-1}) + z_t^\varpi$, with z_t^ϖ accounting for additional factors that influence the real aspirational wage.

The system presented in Subsections 3.1 to 3.4 differs from the framework of Bernanke and Blanchard (2023) along several key dimensions. In terms of time frame, the paper primarily concentrates on medium-run cyclical patterns. However, achieving this requires modeling both short-run and long-run fluctuations.

In the long run, nominal frictions are absent, and macroeconomic variables remain unaffected by shocks from the short and medium run. This is consistent with the classical dichotomy, which suggests that monetary policy does not influence real economic factors in this time frame. In the context of the long run, it is assumed that, in a steady state, real wages and output per employee grow at the same rate. Additionally, an equilibrium relationship between output, employment, and real wages is permitted in the long run, which can deviate from linear trends. Essentially, these variables each follow a unit-root process with a drift.

Central to this paper is the medium run, where nominal frictions come into play and the classical dichotomy breaks down. In the medium run, nominal shocks have an impact on both nominal and real variables. Since medium-run shocks exclude high-frequency disturbances, the medium-run trend offers a clearer representation of the dynamic trajectory of variables. Specifically, when discussing inflation, the concept of “underlying inflation” emerges as a medium-run phenomenon, stripping away short-run volatility to reveal foundational, persistent price movements. Importantly, for monetary policy, the future interest rate path is shaped by the medium-run trajectory of inflation. Lastly, short-run fluctuations are modeled as white noise.

Similar to Bernanke and Blanchard (2023), exponential smoothing is employed to model expectations. In the current framework, it captures medium-run inflation expectations. Long-run expectations are anchored to a policy rate of 0.5 percent, measured at a quarterly frequency. Bernanke and Blanchard (2023) consider both long-run and short-run expectations. While they also revert to the exponential smoothing equation, they directly incorporate forward-looking expectations using measures of expectation taken from the Federal Reserve. In contrast, inflation expectations in the current paper are modeled as a function of past data only. In a high-inflation environment, this approach might have its shortcomings.

Contrary to the Bernanke-Blanchard approach, nominal wages are decoupled from direct price determination in the price Phillips

curve. Meanwhile, inflation expectations directly influence price setting beyond their manifestation through the wage equation. To further allow for price and wage corrections toward their long-run values, the wage gap is incorporated into both the wage and price equations. The wage gap represents the deviation of the real wage from its long-run trend. In contrast, the “catch-up” effect in Bernanke and Blanchard (2023) is modeled as the four-quarter average of CPI inflation minus the one-year inflation expectation from four quarters earlier.

Given the general equilibrium nature of the empirical system in this paper, both the output and employment gaps are endogenous. This contrasts with the Bernanke-Blanchard approach, where domestic slack is treated as exogenous. This distinction is crucial because, in the medium run, output is highly susceptible to changes in inflationary pressures, either directly through demand as modeled in this paper, or indirectly via the monetary policy stance. As a result, the system presented in this paper offers a richer set of dynamics. Lastly, an explicit distinction is made between the output gap and the employment gap, linking them through a modified Okun’s law (see Galí 2015).

In terms of empirical implementation, the current system comprises four observable variables, each with 113 observations. It also includes 18 unobservable endogenous variables and 30 parameters that are estimated. Except for the steady-state inflation, which is fixed at 0.5, all other parameters are estimated. All parameters of the empirical framework are determined using Bayesian methods, and latent variables are extracted with the Kalman filter. In contrast, Bernanke and Blanchard (2023) directly control for all variables; therefore, there are no priors that could influence the estimates and dynamics of the system. Additionally, they allow for a much broader lag structure when estimating their equations.

References

- Adjemian, S., H. Bastani, M. Juillard, F. Karamé, F. Mihoubi, W. Mutschler, J. Pfeifer, M. Ratto, N. Rion, and S. Villemot. 2022. “Dynare: Reference Manual, Version 5.” Dynare Working Paper No. 72, CEPREMAP.

- Benigno, P., and G. Eggertsson. 2023. “It’s Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve.” NBER Working Paper No. 31197.
- Bernanke, B., and O. Blanchard. 2023. “What Caused the U.S. Pandemic-Era Inflation?” Working Paper No. 23-4, Peterson Institute for International Economics.
- Blanchard, O. 2016. “Do DSGE Models Have a Future?” Policy Brief No. 16-11, Peterson Institute for International Economics.
- Bobeica, E., M. Ciccarelli, and I. Vansteenkiste. 2019. “The Link Between Labor Cost and Price Inflation in the Euro Area.” ECB Working Paper No. 2235.
- Calvo, G. 1983. “Staggered Prices in a Utility-Maximizing Framework.” *Journal of Monetary Economics* 12 (3): 383–98.
- Cevik, S., and F. Miryugin. 2023. “Never Different: Fiscal Policy Shocks and Inflation.” IMF Working Paper No. 23/98.
- De Santis, R., and T. Tornese. 2023. “Energy Supply Shocks’ Non-linearities on Output and Prices.” ECB Working Paper No. 2834.
- Eser, F., P. Karadi, P. Lane, L. Moretti, and C. Osbat. 2020. “The Phillips Curve at the ECB.” ECB Working Paper No. 2400.
- Gabaix, X. 2020. “A Behavioral New Keynesian Model.” *American Economic Review* 110 (8): 2271–2327.
- Galí, J. 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Second Edition. Princeton University Press.
- Galstyan, V. 2019. “Inflation and the Current Account in the Euro Area.” Economic Letter No. 4, Central Bank of Ireland.
- . 2021. “Expectations, Unemployment and Inflation: An Empirical Investigation.” *International Finance* 24 (3): 298–312.
- Gelman, A., and Y. Yao. 2021. “Holes in Bayesian Statistics.” *Journal of Physics G: Nuclear and Particle Physics* 48 (1): Article 014002.
- Gertler, M., and J. Leahy. 2008. “A Phillips Curve with an Ss Foundation.” *Journal of Political Economy* 116 (3): 533–72.
- Gonçalves, E., and G. Koester. 2022. “The Role of Demand and Supply in Underlying Inflation — Decomposing HICPX Inflation into Components.” *Economic Bulletin* (European Central Bank) 7/2022.

- Hahn, E. 2020. “The Wage-Price Pass-Through in the Euro Area: Does the Growth Regime Matter?” ECB Working Paper No. 2485.
- Hansen, N.-J., F. Toscani, and J. Zhou. 2023. “Euro Area Inflation after the Pandemic and Energy Shock: Import Prices, Profits and Wages.” IMF Working Paper No. 23/31.
- Harding, M., J. Lindé, and M. Trabandt. 2023. “Understanding Post-COVID Inflation Dynamics.” *Journal of Monetary Economics* 140 (Supplement): S101–S118.
- Hyndman, R., A. Koehler, K. Ord, and R. Snyder. 2008. *Forecasting with Exponential Smoothing: The State Space Approach*. Springer Series in Statistics.
- Ireland, P. 2004. “Technology Shocks in the New Keynesian Model.” *Review of Economics and Statistics* 86 (4): 923–36.
- Iskrev, N. 2010. “Local Identification in DSGE Models.” *Journal of Monetary Economics* 57 (2): 189–202.
- Lane, P. 2022. “Inflation Diagnostics.” *The ECB Blog*, November 25.
- Lo Bello, S., and E. Viviano. 2024. “Some Considerations on the Phillips Curve after the Pandemic.” Occasional Paper No. 842, Banca D’Italia.
- Lorenzoni, G., and I. Werning. 2023. “Wage Price Spirals.” Manuscript, MIT.
- Mankiw, G., and R. Reis. 2002. “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve.” *Quarterly Journal of Economics* 117 (4): 1295–1328.
- Müller, U., and M. Watson. 2015. “Low-Frequency Econometrics.” NBER Working Paper No. 21564.
- Qu, Z., and D. Tkachenko. 2012. “Identification and Frequency Domain Quasi-Maximum Likelihood Estimation of Linearised Dynamic Stochastic General Equilibrium Models.” *Quantitative Economics* 3 (1): 95–132.
- Shapiro, A. 2022. “Decomposing Supply and Demand Driven Inflation.” Working Paper No. 2022-18, Federal Reserve Bank of San Francisco.
- Smets, F., and R. Wouters. 2007. “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach.” *American Economic Review* 97 (3): 586–606.

Werning, I. 2022. "Expectations and the Rate of Inflation." NBER Working Paper No. 30260.

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

What Drives Euro-Area Financial Markets? The Role of U.S. Spillovers and Global Risk*

Lennart Brandt,^a Arthur Saint Guilhem,^b
Maximilian Schröder,^c and Ine Van Robays^a

^aBank of England

^bBanque de France

^cEuropean Central Bank

Understanding euro-area financial market dynamics requires looking beyond borders. We introduce a novel Bayesian structural VAR model designed to jointly estimate the main drivers of euro-area financial asset prices on a daily basis. The model explicitly accounts for transatlantic spillovers and the distinctive influence of the United States, particularly through its monetary policy and its role as a safe haven during periods of heightened risk aversion. The results show that euro-area financial markets are influenced by a complex interplay of domestic factors, U.S. influences, and global risk shocks, which can either amplify or mitigate financial dynamics at home. These significant spillovers can pose substantial challenges for monetary policy, particularly amidst economic and monetary policy divergence.

JEL Codes: C32, C54, E44, E52.

1. Introduction

Monetary policymakers closely monitor daily fluctuations in financial asset prices such as bond yields, stock prices, and exchange rates in search for signals on the underlying expected dynamics of

*We would like to thank Zymantas Budrys, Natalie Burr, Giacomo Carboni, Massimo Ferrari Minesso, Roberto de Santis, Maren Froemel, Catherine L. Mann, and Roberto Motto for fruitful discussions of the model. The views expressed here are those of the authors and do not necessarily reflect those of the European Central Bank, the Eurosystem, the Bank of England, or Norges Bank. Author emails: Lennart.Brandt@bankofengland.co.uk, Arthur.Saint-Guilhem@banque-france.fr, Maximilian.Schroder@ecb.europa.eu, Ine.Van-Robays@ecb.europa.eu.

the economy. Asset prices contain a rich set of information; they react in real time to economic news and hint at how markets assess the effect of recent news or policy decisions on growth and inflation.

At the same time, extracting information from asset prices in real time is challenging; economic shocks do not occur in isolation and can steer asset prices in different directions, blurring their signaling content. Euro-area financial markets are also largely driven by international factors, which adds to the complexity. Two particularly important sources of global impulses have been documented in the literature.

First, many studies have documented the importance of U.S. monetary policy spillovers (Bruno and Shin 2015; Dedola, Rivolta, and Stracca 2017; Buch et al. 2019; Jarociński 2022; Ca’Zorzi et al. 2023), reflecting the central role of U.S. monetary policy in steering the global financial cycle (Miranda-Agrippino and Rey 2020; Gourinchas and Rey 2022) and the dominant role of the U.S. dollar in the international monetary system (Ilzetzki, Reinhart, and Rogoff 2019; Maggiori, Neiman, and Schreger 2019; Boz et al. 2022). U.S. developments affect economic conditions abroad through various channels, including demand, the exchange rate, and financial linkages. As a result, not only U.S. monetary policy shocks but also news about the U.S. economy can influence global financial markets (Albuquerque and Vega 2009; Pinchetti and Szczepaniak 2024).

Second, global financial markets are vulnerable to sudden shifts in global risk sentiment, as shown by the Global Financial Crisis of 2008 and the outbreak of the COVID-19 pandemic in 2020. During such episodes, global investors typically rotate from risky into safe assets amid heightened global risk aversion, which affects the pricing of financial assets globally (Baele et al. 2020; Ahmed 2023). The safe-haven status of U.S. dollar-denominated assets implies that the U.S. dollar and U.S. Treasury yields are particularly susceptible to shifts in global risk sentiment (Habib and Stracca 2012; Fatum and Yamamoto 2016; Bodenstein et al. 2023; Georgiadis, Müller, and Schumann 2024), which in turn can affect asset pricing in the euro area.

To gauge the driving forces behind euro-area financial markets, it is important to jointly account for the effects of these two types of shocks—U.S.-specific shocks and global risk shocks. The academic literature so far has mostly evaluated these in isolation. This

leaves open the question of their relative importance, which is key to understanding the policy implications of a given change in financial conditions. For instance, a decline in euro-area interest rates is likely to exert very different effects depending on whether it reflects a domestic monetary policy easing, negative macroeconomic news, spillovers from U.S. monetary policy accommodation, or heightened global risk aversion.¹ Disentangling these different drivers of financial market movements in a simultaneous way is therefore an important endeavor for monetary policymakers and central bankers at large.

To address this challenge, this paper proposes a parsimonious unified framework that allows separating daily movements of key euro-area financial variables into their underlying drivers, while accounting for the unique role of the U.S. in shaping global financial market dynamics. The primary goal of the model is to provide a real-time evaluation of what steers key euro-area asset prices on a daily basis. We stand out from prior research by explicitly considering international channels² and integrating various driving factors in one common framework,³ which allows us to untangle their relative importance. In addition, our approach departs from the high-frequency literature identifying monetary policy shocks in a narrow time interval around monetary policy announcements (e.g., Jarociński and Karadi 2020), in that we identify euro-area and U.S.

¹The fact that shock dependence matters for the interpretation of changes in asset prices has been documented in the literature on exchange rate pass-through (Forbes, Hjortsoe, and Nenova 2018), for instance, but it applies to financial asset prices more generally.

²For example, Matheson and Stavrev (2014) estimate a daily BVAR model for U.S. yields and stock prices based on sign restrictions exploiting cross-asset price movements. Compared to them, we extend the model to a two-country euro-area/U.S. framework including not only bond yields and equity prices but also the bilateral USD/EUR exchange rate. Additionally, we build on the role of U.S. dollar-denominated assets as safe-haven assets for global investors to identify a new important source of daily market movements, the so-called global risk factor.

³For example, Ca'Zorzi et al. (2023) compare the international transmission of euro-area and U.S. monetary policy shocks relying on “pure” monetary policy shocks which are identified in a narrow time window and purged for central bank information effects following Jarociński and Karadi (2020). Our paper differs by simultaneously identifying macro, monetary policy, and global risk shocks, which provide a more comprehensive picture of the relative importance of monetary policy shocks in driving financial markets.

shocks on a continuous, daily basis, aligning with the frequency at which market participants adjust their perception of the macroeconomic and policy outlook. This broader approach has the benefit of capturing anticipation effects of monetary policy, which play a pivotal role in the transmission of monetary policy to financial conditions (Cao and Foerster 2013). It also allows capturing changes in the perception of market investors on the monetary policy stance which happen on non-meeting days (Jayawickrema and Swanson 2021; Bauer and Swanson 2023). Our identification leads to an interpretation of monetary policy shocks that is broader and more general than in the standard high-frequency literature, in essence capturing the continuous repricing by markets of monetary policy expectations, similar to Cieslak and Pang (2021). Despite these differences, however, we find that our estimated monetary policy shocks comove significantly with estimates from the high-frequency literature on the days when a monetary policy announcement takes place. We also find that realizations of our monetary policy shocks are significantly larger on monetary policy meeting days than on non-meeting days.

In a first step, we set up a Bayesian vector autoregression (BVAR) model at daily frequency that identifies five different drivers of asset price fluctuations: euro-area and U.S. monetary policy, euro-area and U.S. domestic macro news, and global risk. We rely on an identification approach based on sign and magnitude restrictions following Arias, Rubio-Ramírez, and Waggoner (2018), exploiting the information content of daily comovements in risk-free yields, equity prices, and the U.S. dollar-euro exchange rate. The key novelty is that we explicitly allow for simultaneous spillovers between the euro area and the U.S. and account for the potential importance of global risk in driving financial markets on both sides of the Atlantic. This contrasts our approach from the literature on international spillovers of monetary policy shocks that leverages high-frequency information in proxy SVARs (Miranda-Agrippino and Nenova 2022; Degasperis, Hong, and Ricco 2024) or by a combination of high-frequency policy proxies and traditional sign restrictions (Cesa-Bianchi and Sokol 2022). In particular, Cesa-Bianchi and Sokol (2022) disentangle domestic and foreign dynamics by ruling out spillovers from the domestic to the foreign economy with zero restrictions, whereas we allow for simultaneous

international spillovers by imposing softer magnitude restrictions. This offers a complementary perspective on the identification of international spillovers and cross-country shocks. We extensively test the performance of our model by benchmarking its interpretation of daily financial market movements against our knowledge of clear narrative events, and by comparing the estimated structural shocks with other established standards from the literature. In a second step, we use our set of daily shocks as econometric instruments in a monthly proxy structural VAR model, following Stock and Watson (2012), Mertens and Ravn (2013), Gertler and Karadi (2015), and particularly Caldara and Herbst (2019). We use this monthly model to assess how daily financial dynamics carry over to the real economy, and to illustrate why it matters to differentiate between different drivers of financial variables in the first place.

Two sets of results are worth highlighting. First, to grasp the daily shifts in the set of key financial variables that we consider, it is essential to account for the U.S. influence and the impact of global risk. Our results show that spillovers from the U.S. and changes in global risk aversion can explain close to half of the average variation in selected asset prices in the euro area. This large effect of foreign forces might pose challenges for domestic policymakers. At times, these spillover effects might align with the intended direction set by domestic monetary policy. But there might be instances when these spillovers push asset prices in the opposite direction, and policymakers are faced with the decision to counteract these or not. While U.S. spillovers explain a large part of dynamics in euro-area financial markets, euro-area shocks are found less relevant for U.S. financial markets. Instead, shifts in global risk sentiment determine a larger part of the U.S. financial market dynamics, consistent with its safe-haven role.

The second key finding is that daily movements in financial markets provide valuable information about broader, longer-term expectations for economic growth and inflation. Our proxy VAR results show that the way financial markets affect the overall economy can vary significantly based on the combination of factors driving their developments. This implies that for policymakers, breaking down signals from daily financial market indicators is particularly relevant; it provides a crucial signal on expected core economic trends

and insights into how monetary policy is being transmitted to the real economy.

The remainder of the paper proceeds as follows. The next section explains the setup of the daily Bayesian VAR model, the data, and the sign restriction identification scheme, while Subsection 2.5 empirically evaluates the validity of the model setup. Section 3 reviews the main empirical results, focusing on the factors driving euro-area financial markets since the inception of the euro, with particular attention to cross-Atlantic spillovers and global risk shocks. It then assesses the transmission of the identified shocks to the real economy using a proxy VAR model which is described in Section 4. Section 5 concludes.

2. Identifying Daily Drivers of Euro-Area Asset Prices

2.1 *A Daily Two-Country Structural VAR Model*

In order to analyze the drivers of euro-area financial markets, we set up a two-country structural vector autoregression model (SVAR) containing key asset price data for the euro area and the U.S., which we estimate using Bayesian methods. The underlying drivers of daily changes in these euro-area and U.S. asset prices are identified using sign restrictions exploiting cross-asset price movements on the day of the shock. Using these models to analyze interactions at daily frequency offers a number of benefits to circumvent some of the well-documented limitations of shock identification with low-frequency macroeconomic data. In particular, due to the rich data environment, the estimation can rely on rather minimalistic prior assumptions.

Following the notation of Rubio-Ramírez, Waggoner, and Zha (2010), the structural form of the VAR is given by

$$y_t' A_0 = x_t' A_+ + \varepsilon_t'. \quad (1)$$

Here, y_t is an $n \times 1$ vector of endogenous variables and x_t is an $(np+k) \times 1$ vector which collects all predetermined variables, that is, p lags of the endogenous variables and k exogenous and deterministic variables. The conformable matrices A_0 and A_+ capture the contemporaneous comovement between endogenous variables and the

VAR parameters of the predetermined variables, respectively. The corresponding reduced form is given by

$$y'_t = x'_t B + u'_t, \quad (2)$$

where $B = A_+ A_0^{-1}$ and $u_t = \varepsilon'_t A_0^{-1}$ with $E[u_t u'_t] = \Sigma = (A'_0 A_0)^{-1}$. For an orthonormal $n \times n$ rotation matrix, Q , (A_0, A_+) can alternatively be expressed as the observationally equivalent $(A_0 Q, A_+ Q)$. Since the set of possible orthogonal matrices is infinite, there exists an infinite number of structural models that give rise to the same observable data.⁴ However, given a specific Q , there exists a one-to-one mapping between reduced and structural form. Given the Cholesky factorization⁵ of the covariance matrix of the residuals, we can write the reduced form in terms of the underlying structural shocks and the rotation matrix Q (Arias, Rubio-Ramírez, and Waggoner 2018):

$$y'_t = x'_t B + \varepsilon'_t Q' \text{chol}(\Sigma). \quad (3)$$

Since B and Σ are exactly identified and Q uniquely links these to a valid structural model, Q can be used as a parsimonious way to carry theoretical restrictions on the structural parameters that are independent of the reduced form. In particular, these restrictions can be expressed as a Bayesian prior over Q that assigns zero probability mass to parts of the parameter space where the restrictions do not hold. This forms the basis for identification. Statistics, such as impulse response functions or historical decompositions, can be computed in a standard way. In general, however, this identification will not be unique and the imposed sign restrictions will yield a set of models that is plausible under these sign restrictions.

2.2 Priors and Estimation

In the empirical exercises, we follow Arias, Rubio-Ramírez, and Waggoner (2018) and impose a conjugate normal-inverse-Wishart

⁴Because $B = A_+ A_0^{-1} = A_+ Q \cdot Q^{-1} A_0^{-1}$ and $\Sigma = A_0 Q \cdot Q^{-1} A_0^{-1} \forall Q \in \mathcal{O}_n$.

⁵Or any other decomposition h for which it holds that $h(\Sigma)' h(\Sigma) = \Sigma$ (Arias, Rubio-Ramírez, and Waggoner 2018).

prior on the reduced-form parameters as well as a uniform prior on the rotation matrix, Q . The joint prior is hence given by

$$p(B, \Sigma, Q) = P(B, \Sigma) \cdot P(Q). \quad (4)$$

Together with a normal likelihood, this parameterization gives rise to a so-called normal-generalized-normal posterior distribution over the structural parameters (A_0, A_+) conditional on the restrictions imposed. Because $p(B, \Sigma | Y_T)$ is itself a proper posterior density—i.e., $p(B, \Sigma | Y_T) \propto p(B, \Sigma) \cdot p(Y_T | B, \Sigma)$ —and because Q is not informative about the reduced form—i.e., the reduced form is unaffected by the choice of restrictions—this posterior is given by

$$p(B, \Sigma, Q | Y_T) \propto p(Q) \cdot p(B, \Sigma) \cdot p(Y_T | B, \Sigma). \quad (5)$$

In the empirical application, estimation is implemented with Algorithm 1 of Arias, Rubio-Ramírez, and Waggoner (2018), which yields i.i.d. draws from the normal-generalized-normal posterior conditional on the set of sign restrictions. In doing so, we do not restrict the set of (A_0, A_+) itself but the contemporaneous impulse response functions. Each iteration, we draw a naive candidate for the impulse response function which we repeatedly rotate with draws of Q . If the rotated impulse response function conforms to our sign restrictions, we retain the draw; otherwise, we discard it.

2.3 Data

Our set of variables comprises the 10-year euro-area overnight index swap (OIS), the EURO STOXX price index, the U.S. S&P 500, and the bilateral euro-U.S. dollar exchange rate, as well as the spread between the 10-year euro OIS rate and 10-year U.S. Treasury yield.⁶ We use the 10-year euro-area (EA) OIS rate and U.S. Treasury yield, as they are the main references for long-term risk-free rates in both

⁶The yield and spread variables enter the model in first difference, while stock prices and the exchange rate are in log-difference. All variables are at business-day frequency.

jurisdictions.^{7,8} Including long-term rates in the model, as opposed to short-term rates, allows for a broader definition of monetary policy shocks and might better capture shifts in risk sentiment. Regarding monetary policy shocks, using the 10-year yield allows capturing not only conventional monetary policy—which mostly affects the yield curve at short to medium maturities—but also unconventional measures such as those deployed by the Federal Reserve (the Fed) following the Global Financial Crisis. Large-scale asset purchases (LSAPs) have been shown to indeed exert a larger impact at longer maturities by compressing term premia.⁹ In addition, focusing on longer-term risk-free rates is instrumental to identify the global risk shock, as flight-to-safety episodes primarily affect investors' appetite for risk-free bonds against riskier assets, and the U.S. Treasury market is the largest provider of safe, U.S. dollar-denominated assets globally. Cieslak and Povala (2015) and Cieslak and Pang (2021) have shown that risk premium shocks more strongly affect longer maturities of U.S. Treasury bonds, moving 10-year yields more than twice as strongly as 2-year yields. The model is estimated over the period January 1999–December 2023, including four lags of the dependent variables.¹⁰

2.4 Shock Identification

Shock identification is achieved by way of imposing sign restrictions as in Arias, Rubio-Ramírez, and Waggoner (2018). Our approach

⁷The swap rate is the fixed rate which banks engaging in swap contracts agree to pay in exchange for receiving the average overnight interest rate for the duration of the swap. Unlike an unsecured interbank loan, in which the lender is exposed to the full credit risk of the borrower, a swap contract is settled in notional amounts, namely without involving a physical exchange of principals, and thus it is considered near risk-free.

⁸An alternative would be to use the 10-year German bund as a measure of the long-term risk-free rate. However, given that the 10-year OIS and 10-year German bund are tightly correlated over time, estimating the model with the 10-year German bund instead of the 10-year OIS leaves the results largely unaffected.

⁹The downward pressures on the term premium stem from the central bank extracting duration risk from the market, which is the well-documented duration channel of asset purchases. For an empirical assessment of the duration channel of European Central Bank (ECB) asset purchases based on a term structure model, see Eser et al. (2023).

¹⁰While our main interest is on the most recent years, we include the pre-crisis sample as a way to benchmark post-crisis developments vis-à-vis “normal” times.

consists of imposing sign restrictions on daily cross-asset price movements, motivated by the intuition often portrayed in market commentaries that these comovements reveal information about the underlying drivers of financial markets. For instance, bond yields and stock prices typically comove positively in response to news about the economy and negatively in response to news related to monetary policy. This intuition finds strong empirical support (Rigobon and Sack 2003, 2004; Matheson and Stavrev 2014), and has been formalized in several papers that identify the contribution of monetary policy or growth shocks to yield and stock price reactions around monetary policy announcements.¹¹ The imposed sign restrictions are also in line with well-established closed- and open-economy structural model dynamics, such as those described in Bernanke, Gertler, and Gilchrist (1999) and Christiano, Motto, and Rostagno (2014) for the domestic shocks, and Galí and Monacelli (2005), Dedola and Lombardo (2012), and Dedola, Karadi, and Lombardo (2013) for the cross-country spillovers.

A crucial departure from the literature using higher-frequency asset price dynamics for shock identification lies in our broader framework, and this along two dimensions. First, we do not confine the estimates to a specific narrow time window surrounding particular announcements, but instead allow shock identification in “continuous time.” This is consistent with the observation that market-moving news, which influences perceptions on monetary policy or the growth outlook, unfolds virtually every day. Nuances in monetary policy communication which move financial markets often occur outside of these narrow time frames (Jayawickrema and Swanson 2021; Bauer and Swanson 2023), for example. A second departure from this literature is that we identify multiple shocks jointly, which is a key contribution, as it allows us to compare

¹¹For example, Jarociński and Karadi (2020) use high-frequency comovement of interest rates and stock prices around policy announcements to disentangle monetary policy surprises from central bank information shocks. Andrade and Ferroni (2021) use comovements between yields and inflation expectations to disentangle “Odyssean” monetary policy signals from “Delphic” monetary policy surprises (Campbell et al. 2012). Cieslak and Schrimpf (2019) use comovements between stock prices and bond yields at different maturities to decompose asset price fluctuations into monetary policy surprises, information (or growth) surprises, and risk shocks.

Table 1. Sign Restriction Table

	EA Monetary Policy	EA Macro News	U.S Monetary Policy	U.S. Macro News	Global Risk
EA Long-Term Yields	+	+	+	+	−
EA Equity	−	+			−
U.S. Equity			−	+	−
USD/EUR	+	+	−	−	−
EA-U.S. Yield Spread	+	+	−	−	+

Note: Empty fields leave that parameter unrestricted, “+” and “−” denote an increase or a decrease in the respective variable on impact. A “+” (−) for the USD/EUR denotes an appreciation (a depreciation) of the euro vis-à-vis the U.S. dollar. The global risk shock displays a worsening of global risk sentiment.

the relative importance of the key driving factors, while taking a consistent approach to identification.¹²

Table 1 summarizes the structural shock identification scheme via sign restrictions on the contemporaneous impulse response functions. A “+” and “−” denote an increase or decrease, respectively, in the variable following a specific shock, while empty fields leave that response unrestricted. All restrictions are imposed on impact, reflecting the fact that markets typically react to news within the same day.

As shown in Table 1, monetary policy shocks differ from macro shocks through the implied comovement between yields and stock prices; a restrictive monetary policy shock drives yields up and equity prices down, while a positive macro shock increases yields and equity prices (Matheson and Stavrev 2014; Altavilla et al. 2019; Cieslak and Schrimpf 2019; Jarociński and Karadi 2020). Both shocks are assumed to affect the exchange rate similarly, with higher domestic interest rates leading, *ceteris paribus*, to an appreciation of the domestic currency (Engel 2014; Ca’Zorzi et al. 2023). We leave the

¹²For shocks other than the monetary policy shocks, such as the global risk or macro shocks, defining a narrow window around a specific event for identification is more challenging and arguably more arbitrary.

effect of a domestic shock on foreign equity prices unrestricted, as the effect is *a priori* ambiguous.¹³

Second, we allow for instantaneous spillovers between the euro area and the U.S., which requires us to properly control for the country of origin. We address this issue by assuming that a domestic shock has a larger effect on the domestic variable than on the foreign variable. In practice, we apply this magnitude restriction through the spread between euro-area and U.S. yields. For example, a monetary policy tightening in the U.S. might put upward pressures on U.S. and euro-area yields, but as the spread of euro-area over U.S. rates declines, U.S. yields will react more strongly. This contrasts our approach with that of Cesa-Bianchi and Sokol (2022), who rule out direct spillovers from the domestic to the foreign economy by means of stricter zero restrictions.

Finally, a global risk shock causes investors to reallocate their holdings of risky versus safe assets, which causes equity prices and bond yields to move in the same direction. For example, higher risk aversion will reduce appetite for risky assets (Baele et al. 2020) while causing inflows into safe assets. During episodes of heightened global risk, investors flee into U.S. dollar-denominated assets, as these are considered liquid and safe (Farhi and Gabaix 2016; Jiang, Krishnamurthy, and Lustig 2021), which holds in particular for U.S. Treasury bonds in that they embed a convenience yield (Krishnamurthy and Vissing-Jorgensen 2012; Du, Im, and Schreger 2018). Due to the higher liquidity and safety of U.S. Treasuries, we assume that Treasury rates will fall more strongly than euro-area yields upon a worsening in global risk sentiment, which causes the euro-area–U.S. rate spread to decline. In addition, the U.S. dollar is assumed to appreciate vis-à-vis the euro exchange rate, reflecting the role of the U.S. dollar as safe-haven currency (Georgiadis, Müller, and Schumann 2024), and the U.S. economy providing insurance in times of global stress (Gourinchas and Rey 2022).

¹³News about the domestic macroeconomic outlook may drive a reallocation from domestic risky assets into foreign risky assets, which all else equal would exert opposite effects on domestic and foreign equity prices. On the other hand, positive domestic news on the economy can lift foreign equity prices through the demand channel.

Interpretation of the Structural Shocks. Our choice to identify the structural shocks in “continuous time” by relying on comovements in asset prices has implications for the interpretation of the shocks. We identify the structural shocks through the lens of financial markets, with the shocks capturing changes in the interpretation of market participants of daily news on the macroeconomic outlook, global risk, and monetary policy, as captured by changes in asset pricing. This implies that we attain a wider definition of the structural shocks than part of the literature does. For example, compared to the studies using high-frequency monetary policy shock identification during narrow time windows (e.g., Jarociński and Karadi 2020), our monetary policy shocks capture changes in the *perception* of market participants of the monetary policy stance, similar to Cieslak and Pang (2021). These will include shifts in interest rates on days of monetary policy meetings and speeches, but also shifts caused by news that induce market participants to reprice their monetary policy expectations. This could happen following important data releases when monetary policy is data dependent, for example. As such, our monetary policy shocks encompass a broader range of innovations beyond those directly stemming from surprises in monetary policy decisions or communication, as we identify these shocks through asset price movements. Similarly, the global risk shock captures market reactions to events that cause investors to price in increased risk aversion, by shifting toward risk-free assets and safe U.S. dollar-denominated assets. For each day, typically a multitude of shocks is identified.

There are advantages to our identification approach. For example, compared to the literature exploiting market movements over narrow time windows around monetary policy events, there are two main differences. First, our approach is not restricted to days of monetary policy announcements, but we allow for monetary surprises to occur any day. This is consistent with the findings of Jayawickrema and Swanson (2021) and Bauer and Swanson (2023), who have underlined the importance of broadening the time window to a wider set of events in order to accurately pin down monetary policy surprises and central bank information effects. Particularly when monetary policy takes a data-driven approach, markets might reassess the expected stance of monetary policy more frequently. Second, because we do not restrict the time window to specific events, our approach

also allows us to capture in full anticipation effects that materialize ahead of the actual policy announcement through nuances in monetary policy communication on non-monetary policy meeting days such as in speeches. These anticipation effects are an important channel of monetary policy transmission, as shown by Cao and Foerster (2013) who use surveys, newspaper articles, and Internet searches to show that announcements for various rounds of LSAPs in the U.S. were largely expected. Our results, discussed below, confirm that monetary policy announcements are typically largely anticipated by market participants—and our framework effectively categorizes these as monetary policy shocks.¹⁴ Moreover, we find that realizations of monetary policy shocks are significantly larger on monetary policy meeting days than on non-meeting days, which reflects our broader definition of monetary policy shocks.¹⁵

Despite these conceptual differences, we find that our monetary policy shocks comove strongly with estimates from the high-frequency literature on announcement days. To illustrate the validity and information content of the shocks retrieved from our daily SVAR, we project onto them the high-frequency surprises in OIS rates across the term structure, using the changes in OIS rates around ECB announcements from the EA Monetary Policy Database of Altavilla et al. (2019).

For each high-frequency change hf in the OIS rate at maturity τ recorded on announcement day t , we run the following univariate regression, where $\varepsilon^{(i)}$ is the i 'th shock in the daily SVAR:


$$hf_{\tau,t} = \alpha^{(i)} + \beta_{\tau}^{(i)} \varepsilon_t^{(i)} + \nu_{\tau,t}^{(i)}. \quad (6)$$

Figure 1 shows that the shocks from the daily SVAR have significant explanatory power for high-frequency movements in financial market variables around policy announcements. We find that

¹⁴In particular, our results suggest that, once accounting for anticipation effects, there is no evidence that the marginal impact of LSAP programs on long-term interest rates has declined over time. This finding is in line with the evidence provided by Cahill et al. (2013), Cao and Foerster (2013), and Carlson et al. (2020). Cahill et al. (2013) show that, after controlling for pre-announcement market expectations, the yield impact of the Fed's asset purchases has not reduced over time.

¹⁵See Figure A.1 in Appendix A.

Figure 1. Explanatory Power of Euro-Area Shocks for High-Frequency Changes in OIS Rates around ECB Policy Announcements



Note: Panel A shows the term structure of regression coefficients from Equation (6). Error bars denote the 95 percent confidence interval. Panel B shows the associated share of variance explained. Orange lines and grey swathes additionally show the respective parameters of the regression with the first principal component of OIS with maturity up to one year on the left-hand side; as in, for example, Nakamura and Steinsson (2018) and Ca’Zorzi et al. (2023).

our monetary policy shock explains a significant share of the variation in OIS rates around policy announcements—around 40 percent beyond the four-year mark. Since our shocks are identified using a fairly long risk-free rate and estimated on a sample which includes a long period of unconventional monetary policy with short-term interest rates near the effective lower bound, it is expected that they comove most with surprises at medium- to long-term maturities. Point estimates of the parameter β are statistically significant (at the 5 percent level) and sizable with a one-standard-deviation monetary policy shock increasing OIS surprises by about 2.5 basis points (compared to a standard deviation of around 5 basis points for most OIS surprises).

Possible Extensions to Our Framework. While our baseline specification is rather parsimonious, it could be extended to capture additional drivers of euro-area financial conditions, such as sovereign or corporate credit risk. Additionally, one could consider

yields at other maturities as a way to capture different dimensions of monetary policy (e.g., negative interest rate policy versus rate forward guidance versus the ECB's asset purchase program).^{16,17} While changing the specification of the model may enrich the interpretation of monetary policy factors, the results remain largely unaffected.¹⁸ For our global risk shock, this implies that in essence, a simple framework that rests on two types of assets—risk-free bonds and equity—in combination with a safe-haven currency such as the U.S. dollar seems sufficient to pin down the “flight-to-safety” component.

2.5 *Model Interpretation and Validation*

We check the validity of our identification strategy in two ways. First, we inspect the estimated shocks and compare them to alternative empirical measures over time (Section 2.5.1). Second, we review the shock decomposition over a selection of significant market events for which we have a strong prior knowledge on the dominant factor driving market developments over that specific event (Section 2.5.2).

2.5.1 *Economic Interpretation of the Identified Drivers*

Monetary Policy Shocks. An inspection of the estimated euro-area and U.S. monetary policy shocks shows that over periods when the ECB and the Fed have been most active in deploying unconventional policy measures,¹⁹ large realizations of the shocks


¹⁶In practice, empirically disentangling the different monetary policy instruments in one single framework proves challenging. Rostagno et al. (2021) propose an approach that overcomes this challenge based on combining yield curve identification with a large macro BVAR.

¹⁷One important implication of using long-term yields is that the estimated monetary policy shock spans multiple dimensions in line with the target, path, and quantitative easing (QE) factors from the literature. Appendix B provides a cross-check of our estimated monetary policy shock against the target, path, and QE instruments from Miranda-Agrippino and Nenova (2022). The results indicate that our monetary policy shock loads positively on all three factors and can therefore be interpreted as spanning multiple monetary policy dimensions. See Figures B.1–B.4.

¹⁸Appendix B provides robustness checks of alternative model specifications—for example, using sovereign yields.

¹⁹We focus on (i) 2014–20 for the euro area, covering the adoption of negative interest rate policy (NIRP) and the asset purchase program (APP) announcement and the measures taken in response to the COVID-19 pandemic such as the

Figure 2. Euro-Area and U.S. Monetary Policy Shocks during Periods of Unconventional Policy



Note: The charts show the respective posterior mean over all accepted draws from the BVAR.

coincide with important monetary policy announcements (Figure 2). For example, large shocks for the euro area come with the first asset purchase program (APP) announcement in early 2015, its

Pandemic Emergency Purchase Programme (PEPP); and (ii) the period 2008–12 for the U.S., during which the various large-scale asset purchase (LSAP) programs were announced. The choice to focus on this specific period is arbitrary; also beyond this time frame important monetary policy shocks are well captured.

subsequent recalibrations, and the announcement of the Pandemic Emergency Purchase Programme (PEPP) at the onset of the pandemic. For the U.S., the largest negative monetary policy shock during that time frame occurred on the day of the first large-scale asset purchase (LSAP) announcement in November 2008. We also inspect other policy announcements and find a significant contribution of monetary policy shocks in shaping financial conditions on these days. Overall, this indicates that our approach effectively captures key monetary policy announcements.²⁰


Macro Shocks. We assess the euro-area and U.S. macro shocks by comparing them with the Citigroup Economic Surprise Index (CESI), a measure of macroeconomic surprises commonly used by market analysts.²¹ As our identification of the macro shock aims at capturing news that affect the market perception of the macroeconomic outlook, the estimated macro shocks should exhibit some degree of comovement with actual macro surprises. In fact, the comparison shows that both measures strongly comove over the estimated sample for both the euro area and the U.S.²² It also shows that the correlation tends to vary over time. In particular, the identified macro factor disconnects at times from actual macro surprises with some persistence. There are two possible explanations for this. First, while the CESI is an aggregation of all different macro releases, market participants might be inclined to give more attention to specific news at specific periods of time. Second, the disconnect between actual macro surprises and the market perception of the macro outlook as estimated in our model is also consistent with the tendency of markets to over- or under-react to news about the state of the economy in bouts of optimism or pessimism, as in

²⁰These findings are consistent with the literature documenting a strong asset price reaction around large-scale asset purchase announcements. For the euro area, see Andrade et al. (2016) and Altavilla, Carboni, and Motto (2021). For the U.S., see Gagnon et al. (2011) and Vissing-Jorgensen and Krishnamurthy (2011).

²¹The CESI is defined as a weighted historical standard deviation of data surprises (actual releases versus the Bloomberg median survey) and is calculated daily in a rolling three-month window. The weights of the economic indicators are derived from relative high-frequency spot foreign exchange impacts of one-standard-deviation data surprises adjusted to include a time decay feature so as to replicate the contained memory of markets.

²²See Figure A.3 in Appendix A.

Figure 3. Comparison of the Global Risk Factor with Alternative Risk Indicators



Note: To align the volatility of our global risk measure to that of other measures for easier visual comparison, the global risk factor is the six-months moving average of the standardized posterior mean global risk shock from the daily BVAR, while financial market indicators are shown as monthly averages. The yellow line shows the Global Uncertainty Measure of Bobasu, Quaglietti, and Ricci (2024). To note, the measures are not fully comparable, because by construction the global risk shock takes positive and negative values (capturing risk-on and risk-off episodes), while financial market indicators measure volatility and only take positive values, hence the two distinct scales.

Benhabib, Liu, and Wang (2016). Alternatively, it might be that under imperfect information market investors take into account a wider set of economic indicators than pure data releases when forming their expectations about the state of the economy and the central bank’s reaction to it, as suggested by Pastor and Veronesi (2009).

Global Risk Shock. To assess how well we capture shifts in global risk, we compare our estimated shock series to alternative risk measures: the U.S. and euro-area indices of stock market volatility (VIX and VSTOXX, respectively), the euro-area CISS (composite indicator of systemic stress) (Hollo, Kremer, and Lo Duca 2012), and the global economic uncertainty measure of Bobasu, Quaglietti, and Ricci (2024), as plotted in Figure 3. Our global risk factor, which aims at capturing episodes of shifting risk aversion and flight to quality, strongly comoves with these alternative risk measures, exhibiting spikes during episodes of market turmoil such as the 2008 Lehman collapse or at the start of the COVID-19 pandemic in early 2020. We also find evidence that the euro-area sovereign debt crisis was

accompanied by a surge in global risk aversion with investors moving out of euro-denominated assets toward U.S.-denominated assets in an attempt to reduce their exposure to redenomination risk.²³ Our global risk factor—being two-sided—tracks symmetrically both adverse and benign risk sentiment episodes. We find, for example, that at times when the euro-area economy was on a recovery track, such as in the mid-2000s or over 2015–18, the risk environment was especially favorable, which may have further supported the recovery.

2.6 *Event-Study Evidence*

As a second test of validation, we study the model performance more systematically over selected events that caused a market reaction following important news on euro-area monetary policy (such as announcements of APP or negative interest rate policy), U.S. monetary policy (such as the LSAP announcement), or relevant euro-area or U.S. macro releases. Table 2 lists the selected events and our prior expectation of the specific shock in our model that should dominate financial market movements on that day. Figure 4 plots the two-day reaction of euro-area long-term rates for each of the selected events, together with the contributions of the underlying drivers.²⁴ The results for U.S. yields, euro-area stock prices, and the exchange rate are reported in Appendix A.

Two findings are worth highlighting. First, the results show that the model identification is successful at capturing the specific nature of each event. For events that are caused by an important monetary policy announcement, the model tends to identify the monetary policy shock as the dominant driver of the change in euro-area long-term rates over the two-day window around the event (see left


²³This is consistent with evidence of investors' fears of a euro-area breakdown, as captured by rising redenomination risk (De Santis 2019).

²⁴On any given day, releases of macroeconomic data, monetary policy announcements, or any other type of event can generally coincide. The daily VAR decomposes the reduced-form residuals of the system into orthogonal structural shocks, conditional on our set of identifying sign restrictions. The historical decompositions that form the basis for our event study utilize the identified VAR to express all variables and hence daily movements as a sum of current and past shocks, conditional on the dynamics of the system. This is why the contributions are generally nonzero anywhere, even though the underlying shocks are orthogonal.

Table 2. List of Selected Events

Date	Announcement	Type of Shock
Sep. 15, 2008	Lehman Brothers Collapse	Global Risk/Macro
Nov. 25, 2008	LSAP1 Ann.	U.S. Monetary Policy
Mar. 18, 2009	LSAP1 Expansion	U.S. Monetary Policy
May 10, 2010	Greek Program + EFSF + SMP Ann.	EA Monetary Policy/Macro
Nov. 3, 2010	LSAP2 Ann.	U.S. Monetary policy
Jul. 26, 2012	Draghi London Speech: “Whatever It Takes”	EA Monetary Policy/Risk
Jun. 19, 2013	FOMC Meeting; Fed Taper Tantrum	U.S. Monetary Policy
Aug. 22, 2014	Draghi Jackson Hole Speech	EA Monetary Policy
Sep. 4, 2014	10 bps DFR Cut and ABSPP/CBPP Ann.	EA Monetary Policy
Jan. 22, 2015	APP Ann.	EA Monetary Policy
Dec. 3, 2015	10 bps DFR Cut and APP Extension	EA Macro/Global Risk
Jan. 4, 2016	Stock Market Sell-Off	EA Macro/Global Risk
Jun. 23, 2016	Brexit Referendum	EA Macro/Global Risk
Nov. 8, 2016	U.S. Presidential Elections	U.S. Macro/Global Risk
Apr. 23, 2017	First Round French Presidential Elections	EA Macro/Global Risk
Feb. 25, 2020	Intensification COVID-19 Crisis	EA Macro/Global Risk
Mar. 12, 2020	March 2020 Governing Council	Global Risk/Macro
Mar. 18, 2020	PEPP Ann.	EA Monetary Policy
Note: “Ann.” Stands for announcement, “DFR” for deposit facility rate.		

Figure 4. Drivers of the Euro-Area 10-Year
OIS Rate around Selected Events



Note: The chart shows the decomposition of the change on the day of the event and the day after. The results for the variables not shown can be found in Appendix A; see Figures A.4–A.7.

panel in Figure 4). Similar conclusions apply to spillovers around U.S. monetary policy events; U.S. monetary policy shocks tend to be the dominant driver of euro-area long-term rates on these days (see middle panel in Figure 4). Finally, also global risk and macro events are well captured (see right panel). For example, the Brexit referendum is associated with a large decline in the 10-year OIS in the euro area (–20 basis points), which is equally split between two factors—global risk and euro-area macro. This decomposition sheds light on the market perception of the implications of Brexit: increased risk aversion leading to investment flows into safe dollar-denominated assets, combined with a sudden negative reappraisal of the euro-area macro landscape through close trade and financial linkages with the U.K.

Second, the exercise also underlines that—while one shock might dominate developments on a given day—other shocks also matter.

This may reflect two elements. First, news typically encompasses diverse topics. Some of these updates are focused on monetary policy, while others revolve around macro information. For example, the ECB Governing Council meetings tend to coincide with the release of important U.S. macro news such as the preliminary release of U.S. non-farm payroll data. This might blur the signal associated to the specific event that one tries to isolate. Second, the announcement may give rise to a reappraisal along dimensions that are not directly related to the original signal. For instance, euro-area monetary policy events tend to be associated not only with a monetary policy shock but also with a reappraisal of the euro-area macro factor. This result is consistent with the well-documented finding that ECB monetary policy events give rise to two types of monetary policy signals, namely a classic “Delphic” component (signal about a turn in the monetary policy stance) as well as an “Odyssean” component (signal on the central bank’s assessment of the economic outlook); see also Jarociński and Karadi (2020) and Andrade and Ferroni (2021). Our approach allows us to effectively isolate the Delphic component of monetary news. At the same time, our results also suggest an important role of non-monetary news—shifts in macro and global risk—in explaining the daily reaction of asset prices around monetary policy events, which is consistent with the findings from the literature (Nakamura and Steinsson 2018; Cieslak and Schrimpf 2019).

3. Empirical Results

3.1 *What Drives Euro-Area and U.S. Financial Markets?*


To get a broader overview of the drivers of euro-area asset prices, we present the historical decomposition of euro-area long-term rates, equity prices, and the euro-dollar exchange rate over three subperiods: (i) January 2008–December 2013, dominated by the Global Financial Crisis and the euro-area sovereign debt crisis (“double-dip recession”), (ii) January 2014–December 2019, containing the shift of the ECB’s monetary policy into unconventional domain, and (iii) the COVID-19 pandemic and the inflation surge of 2022.²⁵

²⁵The corresponding results for the U.S. are shown in Appendix A; see Figure A.8.

Double-Dip Recession. During the Global Financial Crisis, the deterioration of the euro-area domestic macro outlook explained about half of the decline in euro-area risk-free rates and caused close to 20 percent fall in euro-area equity prices compared to early 2008; see the top panel in Figure 5. The ECB's monetary policy response (as captured in the euro-area policy factor) lowered long-term yields while supporting domestic equity markets. The euro fell substantially against the U.S. dollar in the first months of the crisis, reflecting safe-haven flows into dollar-denominated assets (as captured by the global risk shock). As the macroeconomic outlook deteriorated on both sides of the Atlantic, macro shocks did not steer the bilateral exchange rate in a specific direction. That changed in the years thereafter, however, when the U.S. economy recovered from the Global Financial Crisis while adverse macro shocks continued to weigh on financial markets in the euro area as the sovereign debt crisis intensified, lowering long-term yields and equity prices further. The euro exchange rate weakened substantially as investors questioned the robustness of the common currency (captured by a combination of euro-area macro and global risk shocks). These factors more than offset upward pressures on the euro from the Fed's unconventional policy measures in 2011–12 that depreciated the U.S. dollar. The persistent drag of adverse macro risk shocks on the outlook for economic activity and inflation in the euro area triggered the ECB to resort to unconventional monetary policy in the months thereafter.

Unconventional Monetary Policy. Over the period 2014–19, the ECB widened its toolkit to include unconventional measures such as negative interest rates, rate forward guidance, and asset purchases, with the objective to bring inflation back to target in the face of persistent deflationary pressures. Model results show that these measures left a significant footprint on euro-area financial markets. This is particularly visible in the contribution of the policy factor, which lowered euro-area long-term yields by about 60 basis points in cumulated terms over 2014–19. This supported domestic equity markets and significantly depreciated the euro exchange rate throughout the review period; see panel B in Figure 5. The deterioration in the euro-area macro outlook added to the downward pressure on long-term yields and the euro exchange rate. By contrast, the improvement in the U.S. macroeconomic outlook pushed up euro-area long-term yields and stock prices, while adding

Figure 5. Historical Decomposition of Euro-Area Asset Prices



Note: Contributions are rebased to zero at the beginning of the review period. For equity prices and the exchange rate they are rescaled to match cumulative growth in the observable variable in percent instead of log-differences. The corresponding decomposition of the U.S. variables can be found in Appendix A ; see Figure A.8.

downward pressures on the euro exchange rate. Although the Fed started to raise policy rates in December 2015, the U.S. policy factor continued to exert downward pressures on U.S. and euro-area yields, indicating a persistently accommodative stance.

COVID-19 Pandemic and the Inflation Surge. The outbreak of the COVID-19 pandemic in early 2020 affected euro-area financial markets differently. The shock mainly affected markets through inflows into safe assets, lowering euro-area long-term yields further and explaining a substantial part of the fall in equity prices in the first months of 2020. Shifts in global risk sentiment were even more impactful in the U.S.—aligning with its recognized status as a safe haven—and accounted for a significant portion of the strengthening of the U.S. dollar in the early part of the pandemic; see panel C in Figure 5. Domestic macro and monetary policy factors did not point to a clear direction initially, which might be consistent with the highly accommodative ECB monetary policy stance before the pandemic hit.²⁶ Central bank communication might have also played a role. Model results suggest that initial ECB communication failed to fully reassure markets, as visible in the large tightening shock at the early March press conference. This was, however, soon followed by a sharp reversal in the accommodative direction after the March 18 decision to launch PEPP.

In the years thereafter, yields and equity prices increased again as the global economy recovered from the pandemic slump. Between 2021 and 2023, the model attributes a sizable part of the recovery in yields and equity prices to improving expectations on the macro outlook on both sides of the Atlantic. As macro dynamics were stronger in the U.S., however, the euro depreciated against the dollar. At the same time, monetary policy started to tighten at an unprecedented pace as inflation hit record high levels due to the recovery in demand and surging energy prices after Russia invaded Ukraine.²⁷

²⁶With policy rates in negative territory and a sizable balance sheet, markets might have anticipated that room for monetary policy to maneuver was limited, at least regarding the ability to further lower the risk-free yield curve.

²⁷This additional tightening impulse—over and above the reaction to an improving macro outlook—might in part reflect the resurgence of adverse supply factors. In theory, the impact of a supply shock on asset prices is ambiguous and depends on how markets expect the central bank to react. If the shock is perceived to be temporary, markets may expect the central bank to look through,

The simultaneous economic rebound in both the euro area and the U.S., accompanied by a synchronous tightening of monetary policies, largely accounted for the sharp rise in yields in the euro area, while intensifying fluctuations in euro-area equities; see panel C in Figure 5. By contrast, because both euro-area and U.S. factors were of similar nature but tended to offset each other in terms of their exchange rate impact, movements in the euro-U.S. dollar pair were relatively more contained.


Overall, these events illustrate that financial markets in the euro area are influenced by a complex interplay of shocks that do not occur in isolation and can amplify, or suppress, financial dynamics. This is particularly the case for U.S. spillovers to euro-area markets, as illustrated next.

3.2 Cross-Atlantic Financial Spillovers

Financial market dynamics—particularly in the euro area—are significantly determined by cross-Atlantic spillovers. The forecast error variance decomposition indicates that U.S. shocks are found to explain close to 40 percent of variability in euro-area yields and equity prices over the full sample; see Figure 6. Shocks originating in the euro area, while relatively less important, still exert sizable influence on U.S. financial markets, and account for about 30 percent of the dynamics in U.S. equity prices on average. Interestingly, our finding that spillovers might run both ways is similar to Kearns, Schrimpf, and Xia (2023), who focus on monetary policy shocks and find that Fed spillovers tend to dominate ECB spillovers but that ECB spillovers to U.S. financial conditions remain tangible. Euro-area financial markets are also affected by shifts in global risk sentiment, although these types of shocks play a larger role for U.S. financial markets given the safe-haven status of U.S. dollar-denominated assets. Taken together, spillovers from U.S. shocks and

leaving financial variables unaffected. On the other hand, the types of supply shocks observed during the pandemic recovery—especially the sharp increase in energy prices—are likely to have contributed to a more forceful monetary policy response than during “normal” times, hence leading the central bank to deviate from its standard reaction function and adding upward pressures on yields as shown by the model results.

Figure 6. Forecast Error Variance Decomposition



Note: The figure shows the one-step-ahead forecast error variance decomposition, averaged over all accepted draws from the daily BVAR.

changes in global risk aversion can explain close to half of the average variation in selected asset prices in the euro area.

Interestingly, the extent to which foreign shocks drive euro-area financial markets differs across assets. Foreign forces explain more than half of euro-area equity prices over our sample, while this share is less for euro-area yields, possibly reflecting the role of ECB’s rate forward guidance in shielding euro-area shorter-term rate expectations from foreign influence. For the euro-U.S. dollar exchange rate, euro-area-specific shocks only account for about 20 percent of the variability, with U.S. shocks explaining about double that amount. Variation in global risk sentiment is, however, the dominant driver of the bilateral euro-U.S. dollar pair, which is in line with the literature that underscores the importance of risk premia in driving exchange rates (Engel 2014), as well as the safe-haven role of the U.S. dollar (Georgiadis, Müller, and Schumann 2024).


To illustrate the importance of cross-Atlantic financial spillovers, it is useful to focus on periods that brought about major shifts in financial markets. As the extent of spillovers depends on the underlying dominant shock, we study the spillovers during (i) the U.S. dot-com bubble (U.S. macro shock), (ii) the euro-area sovereign debt crisis (euro-area macro shock), (iii) the aftermath of the first round of LSAPs by the Fed (U.S. monetary policy shock), and (iv) the period since the ECB started using unconventional monetary policy (euro-area monetary policy shocks).

The U.S. dot-com bubble period illustrates that U.S.-driven macro shocks can significantly steer euro-area financial markets. After a prolonged period of optimistic macro prospects, the burst of the dot-com bubble in early March 2000 triggered a sharp sell-off in equity markets in the U.S. that largely transmitted to EA equity prices. Adverse U.S. macro shocks caused EA equity prices to decline by about 20 percent over a period of almost two years, which is close to the magnitude of the domestic U.S. equity price reaction. The burst of the bubble also motivated investors to shift their portfolios from equities to bonds—as captured by the global risk shock—which further weighed on euro-area equity prices. The negative U.S. shock came at a time when the ECB was tightening its policy, causing the overall decline in equity prices in the euro area to be much more pronounced than the one observed in the U.S. over that period.

In comparison, the spillovers of euro-area macro shocks to U.S. financial markets are found to be noticeably smaller. Whereas the transmission following the U.S. dot-com bubble across equity markets was close to full, only less than half of the drop in euro-area equity prices caused by adverse macro shocks spilled over to U.S. equity markets following the sovereign debt crisis (see panel B in Figure 7). However, as the debt crisis also triggered flows into safer assets such as bonds, particularly dollar-denominated ones, U.S. equity prices and long-term yields were importantly affected through the global risk channel instead.

Regarding monetary policy spillovers, it is also the case that U.S. monetary policy shocks spill over more significantly to euro-area financial markets than the other way around, which is in line with the literature that underlines the dominant role of U.S. monetary policy in steering the global financial cycle (Miranda-Agrippino and Rey 2020; Ca’Zorzi et al. 2023). After the Fed embarked upon quantitative easing policies in late 2008, accommodative U.S. monetary policy shocks are found to have considerably lowered euro-area long-term yields; see panel C in Figure 7. Overall, as the ECB was loosening its policy stance as well over most of that period, the spillovers from the Fed supported domestic efforts to ease financial conditions. In contrast, the spillovers of the ECB’s unconventional policies since mid-2013—comprising forward guidance, negative short-term rates, and asset purchases—are hardly registered in U.S. long-term yields,

Figure 7. Cross-Atlantic Spillovers in Financial Markets



Note: Panel A shows cumulative changes in EA equity prices from March 10, 2000 to December 31, 2002 and panel B shows cumulative changes in U.S. equity prices from July 7, 2011 to December 31, 2012. Panel C shows cumulative changes in EA 10-year OIS from November 24, 2008 to December 31, 2019, and panel D shows cumulative changes in 10-year U.S. Treasury yields from July 4, 2013 to December 31, 2019. All shock contributions are normalized to zero at the beginning of the review period. The corresponding decomposition of variables not shown here can be found in Appendix A; see Figures A.9–A.12.

as shown in panel D. At that time, rates in the U.S. were partly anchored by rate forward guidance, which might have been a reason why euro-area-to-U.S. spillovers were smaller than those of the Fed's policies to euro-area yields before 2013 when the ECB introduced rate forward guidance as well.

In sum, these findings underscore the necessity of considering external influences when dissecting financial conditions. For the euro area, changes in U.S. monetary policy as well as global risk perceptions play important roles, affecting yields and equity prices, and determining the strength of the euro exchange rate. For U.S. financial markets, by contrast, euro-area-specific dynamics can matter but are typically less relevant than domestic or global risk factors. Overall, looking beyond borders appears key for a proper understanding of financial market dynamics in real time.

4. Transmission to the Macroeconomy: Proxy SVAR

As a final exercise, this section analyzes the response of macroeconomic variables to the shocks extracted in the previous section. In doing so, we draw on the proxy-SVAR literature (see, e.g., Stock and Watson 2012; Mertens and Ravn 2013; Gertler and Karadi 2015) and employ the Bayesian proxy-SVAR algorithm proposed by Caldara and Herbst (2019). The starting point is a VAR at monthly frequency that includes the financial variables contained in the daily model as well as industrial production, inflation, and the unemployment rate for both the euro area and the U.S. Stacking these variables in a vector denoted by Y_t yields the structural form VAR

$$B_0 Y_t = B_+ Y_{t-1} + \varepsilon_t, \quad (7)$$

where B_0 and B_+ contain the structural parameters, and $\varepsilon_t \sim N(0, I)$ is the vector of structural shocks. The reduced-form residuals $u_t \sim N(0, \Sigma)$ are hence given by $u_t = B_0^{-1} \varepsilon_t$, with $\Sigma = (B_0' B_0)^{-1}$ and the reduced-form VAR coefficients by $\bar{B}_+ = B_0^{-1} B_+$, where without loss of generality $\bar{B}_+$ and B_+ collect constants and coefficients on lags.

As with any standard VAR, this model requires additional restrictions for identification due to the many combinations of structural parameters that yield the same reduced-form parameters and hence likelihood function. We leverage this identification challenge to connect high-frequency shocks derived from our daily financial VAR to the set of monthly macroeconomic variables.

In doing so, we draw from the proxy-SVAR literature (see, e.g., Stock and Watson 2012; Mertens and Ravn 2013; Gertler and Karadi

2015) and identify the structural shocks in the monthly VAR one at a time using a single instrument, Z_t , for each economic shock of interest, ε_t^p , and its corresponding row in the inverse impact matrix, B_0^p . In our setup, the natural choice for the instruments are then the structural shocks from our daily BVAR. Formally, we follow Caldara and Herbst (2019) and establish a link between the structural shock, ε_t^p , and the instrument, Z_t , through

$$Z_t = \beta \varepsilon_t^p + \sigma_\nu \nu_t, \quad \text{with } \nu_t \sim N(0, 1) \text{ and } \nu_t \perp \varepsilon_t. \quad (8)$$

Two critical assumptions in the proxy-SVAR literature are directly implied by this mapping:

$$E[Z_t \varepsilon_t^{p'}] = \theta \quad (9)$$

$$E[Z_t \varepsilon_t^{q'}] = 0. \quad (10)$$

Assumption (9) requires that the instrument is correlated with the policy shock of interest, ε_t^p . The relevance of the proxy is directly related to the signal-to-noise ratio β/σ_ν . The higher this ratio, the more information the proxy contains on the structural shock and the closer the time series of the structural shock to the time series of the proxy. Conversely, when $\beta/\sigma_\nu = 0$, then the proxy is driven by noise entirely and does not contain information on ε_t . Weak identification results when β/σ_ν is close to zero. Because the structural shocks extracted from the daily VAR are identified using sign restrictions that embody the economic interpretation of ε_t^p , condition (9) is by construction fulfilled. In addition, prior beliefs about the relationship between the proxy and the structural shock of interest can be directly introduced through the tightness of the prior on σ_ν (Caldara and Herbst 2019). In this exercise, we focus on macro-level responses to our daily financial shocks. Therefore, we favor structural shocks that closely align with the proxy by adopting the high-relevance prior outlined in Caldara and Herbst (2019). This prior suggests that only half of the variation in the proxy can be attributed to measurement error, although our empirical results remain robust even when imposing an uninformative prior.

Assumption (10) requires that the instrument is orthogonal to the other structural shocks, ε_t^q , in the macro VAR. To ensure that this condition continues to hold once we aggregate the daily

structural shocks to monthly frequency, we control for possible correlations during the aggregation process.²⁸

4.1 Proxy SVAR Methodology and Estimation

To estimate the proxy SVAR described above, we employ the algorithm in Caldara and Herbst (2019) throughout, which has key advantages over other algorithms in the literature. Taking the proxy into account, the structural VAR can be augmented such that

$$\tilde{B}_0 \tilde{Y}_t = \tilde{B}_j \tilde{Y}_{t-1} + \tilde{\varepsilon}_t, \quad (11)$$

where $\tilde{Y}_t = [Y_t' Z_t']'$ and $\tilde{\varepsilon}_t = [\varepsilon_t' \nu_t']'$. $\tilde{B}_0$ and $\tilde{B}_j$ contain the SVAR parameters as well as the parameters from proxy equation (8) with

$$\tilde{B}_0 = \begin{bmatrix} B_0 & \mathbf{0} \\ -\frac{\beta}{\sigma_\nu} B_0^p & \frac{1}{\sigma_\nu} \end{bmatrix} \text{ and } \tilde{B}_+ = \begin{bmatrix} B_+ \\ -\frac{\beta}{\sigma_\nu} B_+^p \end{bmatrix}. \quad (12)$$

This gives rise to the likelihood function

$$p(Y_{1:T}, Z_{1:T} | \tilde{B}_0, \tilde{B}_+) = p(Y_{1:T} | \Sigma, \tilde{B}_+) p(Z_{1:T} | Y_{1:T}, B_0, B_+, \beta, \sigma_\nu), \quad (13)$$

which contains not only the likelihood function of the reduced-form VAR but also the likelihood of the proxy given the macroeconomic data and structural parameters. This allows for efficient use of information contained in the proxy, as it is utilized during the estimation of both the structural and reduced-form parameters, while also allowing for the joint modeling of all sources of uncertainty (Caldara and Herbst 2019). In contrast, traditional proxy-SVAR approaches rely on a two-step procedure which does not leverage proxy information during the estimation of the reduced-form VAR.²⁹

²⁸To obtain the monthly instruments, we aggregate the daily structural shock series from the BVAR model over the number of successful draws and over time, which introduces some correlation between the instruments in monthly frequency—even though shocks from each draw are uncorrelated at the daily frequency. To address this issue, we purge the instruments by sequentially regressing these on the other shock series to ensure the shock instruments are uncorrelated.

²⁹The first-stage regression isolates the variation of the reduced-form residual, u_t^p , that can be explained by the proxy. In the second-stage regression, the reduced-form residuals of the other variables in the VAR are regressed on the fitted values, which then allows to compute B_0^p based on the restrictions implied by these regression estimates.

For comprehensive explanations and detailed derivations of the algorithm, readers are directed to Caldara and Herbst (2019).


4.2 Empirical Results: Shock Dependence Matters

Figure 8 shows the estimated macroeconomic effects of an impact increase in long-term rates in the euro area by 10 basis points depending on the shock causing higher yields: (i) a euro-area monetary policy shock, (ii) a U.S. monetary policy shock, (iii) a euro-area macro shock, (iv) a U.S. macro shock, or (v) a global risk shock. Here we report results using the euro-area yields to normalize the impact of the shocks, but the macro implications are the same when another endogenous variable is used to normalize, such as the exchange rate.

The results show that, indeed, the implications of higher yields for euro-area industrial production, inflation, and unemployment differ depending on the underlying driver. For example, a monetary policy shock originating within the euro area leads to higher unemployment and lower output, whereas the same shock originating in the U.S. has largely insignificant effects. Likewise, euro-area inflation responds less strongly to a U.S. monetary policy tightening, likely due to the conflicting effect of the euro depreciation. However, in both cases, the responses are generally insignificant and accompanied by wide confidence bounds. In contrast, when long-term euro-area yields rise due to an improving domestic macroeconomic outlook, industrial production and inflation tend to benefit from stronger demand rather than being dampened by higher yields and an appreciating euro. Similarly, when euro-area yields increase in response to positive U.S. macroeconomic news that spills over into euro-area financial markets, domestic industrial production and inflation increase. Finally, when euro-area risk-free rates are higher because of improving global risk sentiment—in the recovery phase from the Global Financial Crisis or the COVID-19 pandemic, for example—higher yields will also be associated with strengthening growth and rising inflationary pressures.

These findings confirm that looking at daily financial market dynamics is useful for policymakers; their comovements convey valuable information on how changes in financial assets will transmit to the real economy and inflation. Depending on the dominant driver, the implications can differ greatly. This underscores the need

Figure 8. PSVAR Impulse Response Functions for Euro-Area Macrovariables



Note: Responses to the policy and macro shocks are normalized to a shock which increases the respective long-term interest rate by 10 basis points on impact. Responses to the global risk shock are normalized to a shock which causes an appreciation of the bilateral euro-dollar exchange rate by 1 percent on impact. The responses in the U.S. and the financial block are shown in Appendix Figures A.13 and A.14, respectively.


to understand the drivers of financial market dynamics in the first place.

5. Conclusions

Understanding financial market dynamics in the euro area requires looking beyond borders. In this paper, we provide a novel, parsimonious model which jointly estimates key drivers of euro-area financial asset prices. The model explicitly accounts for the unique role of the U.S. in shaping global financial market dynamics and cross-Atlantic spillovers. The results show that euro-area financial markets are driven by a complex interplay of domestic and foreign shocks that can amplify or suppress financial dynamics at home. Spillovers from the U.S. and changes in global risk aversion account for close to half of the average variation in selected asset prices in the euro area. This high degree of spillovers can pose challenges for monetary policymaking and at times counteract domestic efforts to stabilize inflation. Moreover, depending on the driving force, the same change in financial market variables can have different implications for macroeconomic variables. Therefore, a comprehensive understanding of the underlying drivers of financial market dynamics is crucial.


Appendix A. Additional Figures

Figure A.1. Posterior Distribution of Policy Shock Volatilities




Note: The two panels show the posterior distribution of standard deviations of the two policy shocks summed over (i) all overlapping two-day windows, (ii) days of policy meetings and the day after, and (iii) those policy meetings at which rates were changed. The one-day standard deviation of the shocks in the model is standardized to unity. Therefore, the expected standard deviation of two-day sums is roughly $\sqrt{2} = 1.41$.

Figure A.2. Anticipation of ECB and Fed Asset Purchases




Note: Panel A shows the cumulative contribution of the euro-area monetary policy shock to changes in the euro-area 10-year OIS rate from August 1, 2014 to March 31, 2015. Panel B shows that of the U.S. monetary policy shock to the U.S. 10-year Treasury yield from June 3, 2010 to December 31, 2010.

Figure A.3. Comparison of Estimated Macro Shocks with the Citigroup Economic Surprise Index




Note: The charts show the standardized rolling three-month sum of the respective posterior mean over all accepted draws from the daily BVAR.

Figure A.4. Drivers of the Euro-Area Equity Prices around Selected Events




Note: The chart shows the decomposition of the change on the day of the event and the day after.

Figure A.5. Drivers of the Euro-Dollar Exchange Rate around Selected Events




Note: The chart shows the decomposition of the change on the day of the event and the day after.

Figure A.6. Drivers of the U.S. Equity Prices around Selected Events




Note: The chart shows the decomposition of the change on the day of the event and the day after.

Figure A.7. Drivers of the U.S. Treasury Yields around Selected Events




Note: The chart shows the decomposition of the change on the day of the event and the day after.

Figure A.8. Historical Decomposition of U.S. Asset Prices




Note: Contributions are rebased to zero at the beginning of the review period. For equity prices and the exchange rate, they are rescaled to match cumulative growth in the observable variable in percent instead of log-differences.

Figure A.9. Cross-Atlantic Spillovers in Financial Markets: U.S. Dot-Com Crash 2000–02




Note: The charts show the average historical decomposition of changes from March 10, 2000 to December 31, 2002 in asset prices not displayed in Figure 7. All shock contributions are normalized to zero at the beginning of the review period.

Figure A.10. Cross-Atlantic Spillovers in Financial Markets: EA Sovereign Debt Crisis 2011–12




Note: The charts show the average historical decomposition of changes from July 7, 2011 to December 31, 2012 in asset prices not displayed in Figure 7. All shock contributions are normalized to zero at the beginning of the review period.

Figure A.11. Cross-Atlantic Spillovers in Financial Markets: Fed Unconventional Policy 2008–19




Note: The charts show the average historical decomposition of changes from November 24, 2008 to December 31, 2019 in asset prices not displayed in Figure 7. All shock contributions are normalized to zero at the beginning of the review period.

Figure A.12. Cross-Atlantic Spillovers in Financial Markets: ECB Unconventional Policy 2013–19




Note: The charts show the average historical decomposition of changes from July 4, 2013 to December 31, 2019 in asset prices not displayed in Figure 7. All shock contributions are normalized to zero at the beginning of the review period.

Figure A.13. PSVAR Impulse Response Functions for U.S. Macrovariables



Note: Responses to the policy and macro shocks are normalized to a shock which increases the respective long-term interest rate by 10 basis points on impact. Responses to the global risk shock are normalized to a shock which causes an appreciation of the bilateral euro-dollar exchange rate by 1 percent on impact.

Figure A.14. PSVAR Impulse Response Functions for Financial Variables



Note: Responses to the policy and macro shocks are normalized to a shock which increases the respective long-term interest rate by 10 basis points on impact. Responses to the global risk shock are normalized to a shock which causes an appreciation of the bilateral euro-dollar exchange rate by 1 percent on impact.

Appendix B. Additional Robustness Exercises

B.1 Choice of Maturity

Table B.1 presents the results of an alternative estimation of the model using the 2-year rate instead of the 10-year rate. The maturity affects the identification of the various shocks, but overall the results remain close with a highly positive correlation across the shocks from the two specifications.

Table B1. Correlation of Shocks
Relative to Benchmark Specification

	EA Policy	EA Macro	U.S. Policy	U.S. Macro	Global Risk
Two-Year	0.79	0.74	0.85	0.95	0.86

B.2 Alternative Risk Indicators

In Table B.2, we test the robustness of the model by (i) replacing the bilateral euro-U.S. dollar exchange rate with the broad U.S. dollar exchange rate and (ii) replacing the euro-area 10-year OIS rate with the euro-area GDP-weighted sovereign yield (also in the interest rate spread). It appears that the estimated global risk shock in both cases is highly correlated with the global risk shock of our benchmark model when making these changes (correlation of 0.7 and 0.99, respectively), which indicates that the results are robust to these alternative risk indicators.

Table B2. Correlation of Shocks
Relative to Benchmark Specification


	EA Policy	EA Macro	U.S. Policy	U.S. Macro	Global Risk
EA Sov.	0.75	0.74	0.98	0.96	0.99
U.S. NEER	0.92	0.99	0.91	0.74	0.70

However, it is important to note that using euro-area sovereign yields to identify shocks might inadvertently incorporate euro-area-specific risk factors into broader shock definitions, such as the global risk shock and the euro-area monetary policy shock. The correlation of shocks shown in Table B.2 above suggests, for example, that the identified euro-area monetary policy shocks differ when using sovereign yields. This discrepancy may arise because sovereign yields could misinterpret shifts in euro-area sovereign spreads—such as those observed during the euro-area sovereign debt crisis—as monetary policy shocks, which might be problematic.

B.3 Composition of the Monetary Policy Shocks


In this section, we present a robustness exercise where we regress the three monetary policy factors from Miranda-Agrippino and Nenova (2022) (target, path, QE) onto our own monetary policy factor. Several results are worth highlighting. First, all three monetary policy factors load positively on our monetary policy factor with a reasonable fit (see Figure B.1). Second, as expected, our monetary policy shock loads more positively on the path and QE dimensions than on the target dimension. And when using the two-year rate instead, the loading on the target factor increases significantly (Figure B.2). Last, using the target, path, and QE shocks purged from the information effects, as has become usual practice in the literature, tends to improve the fit of the regression (see light blue bars compared to dark blue bars). This is reassuring because our identification scheme to separate monetary policy and macro shocks based on the comovement of yields and stock prices bears similarity with the “poor man’s restriction” of Jarociński and Karadi (2020), which is precisely intended to extract monetary policy shocks distinct from information effects.

Figure B.1. Euro Area: Benchmark Specification




Note: This figure shows coefficients and goodness-of-fit statistics from a regression of the EA monetary policy factors of Miranda-Agrippino and Nenova (2022) onto the mean EA policy shock from the benchmark specification of the daily VAR.

Figure B.2. Euro Area: Alternative Specification




Note: This figure shows coefficients and goodness-of-fit statistics from a regression of the EA monetary policy factors of Miranda-Agrippino and Nenova (2022) onto the mean EA policy shock from an alternative specification using two-year yields in the daily VAR.

Figure B.3. U.S.: Benchmark Specification



Note: This figure shows coefficients and goodness-of-fit statistics from a regression of the U.S. monetary policy factors of Miranda-Agrippino and Nenova (2022) onto the mean U.S. policy shock from the benchmark specification of the daily VAR.

Figure B.4. U.S. Alternative Specification



Note: This figure shows coefficients and goodness-of-fit statistics from a regression of the U.S. monetary policy factors of Miranda-Agrippino and Nenova (2022) onto the mean U.S. policy shock from an alternative specification using two-year yields in the daily VAR.

References

- Ahmed, R. 2023. "Flights-to-Safety and Macroeconomic Adjustment in Emerging Markets: The Role of US Monetary Policy." *Journal of International Money and Finance* 133 (May): Article 102827.
- Albuquerque, R., and C. Vega. 2009. "Economic News and International Stock Market Co-movement." *Review of Finance* 13 (3): 401–65.
- Altavilla, C., L. Brugnolini, R. S. Gürkaynak, R. Motto, and G. Ragusa. 2019. "Measuring Euro Area Monetary Policy." *Journal of Monetary Economics* 108 (December): 162–79.
- Altavilla, C., G. Carboni, and R. Motto. 2021. "Asset Purchase Programs and Financial Markets: Lessons from the Euro Area." *International Journal of Central Banking* 17 (4): 1–48.
- Andrade, P., J. Breckenfelder, F. De Fiore, P. Karadi, and O. Tristani. 2016. "The ECB's Asset Purchase Programme: An Early Assessment." Working Paper No. 1956, European Central Bank.
- Andrade, P., and F. Ferroni. 2021. "Delphic and Odyssean Monetary Policy Shocks: Evidence from the Euro Area." *Journal of Monetary Economics* 117 (January): 816–32.
- Arias, J. E., J. F. Rubio-Ramírez, and D. F. Waggoner. 2018. "Inference Based on Structural Vector Autoregressions Identified with Sign and Zero Restrictions: Theory and Applications." *Econometrica* 86 (2): 685–720.
- Baele, L., G. Bekaert, K. Inghelbrecht, and M. Wei. 2020. "Flights to Safety." *Review of Financial Studies* 33 (2): 689–746.
- Bauer, M. D., and E. T. Swanson. 2023. "A Reassessment of Monetary Policy Surprises and High-frequency Identification." *NBER Macroeconomics Annual 2022*, Vol. 37, ed. M. Eichenbaum, E. Hurst, and V. A. Ramey, 87–155. University of Chicago Press.
- Benhabib, J., X. Liu, and P. Wang. 2016. "Sentiment, Financial Markets, and Macroeconomic Fluctuations." *Journal of Financial Economics* 120 (2): 420–43.
- Bernanke, B. S., M. Gertler, and S. Gilchrist. 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." *Handbook of Macroeconomics*, Vol. 1, Part C, ed. J. B. Taylor and M. Woodford, 1341–93. North-Holland.

- Bobasu, A., L. Quaglietti, and M. Ricci. 2024. "Tracking Global Economic Uncertainty: Implications for the Euro Area." *IMF Economic Review* 72 (2): 820–57.
- Bodenstein, M., P. Cuba-Borda, N. Gornemann, I. Presno, A. Prestipino, A. Queralto, and A. Raffo. 2023. "Global Flight to Safety, Business Cycles, and the Dollar." International Finance Discussion Paper No. 1381, Board of Governors of the Federal Reserve System.
- Boz, E., C. Casas, G. Georgiadis, G. Gopinath, H. Le Mezo, A. Mehl, and T. Nguyen. 2022. "Patterns of Invoicing Currency in Global Trade: New Evidence." *Journal of International Economics* 136 (May): Article 103604.
- Bruno, V., and H. S. Shin. 2015. "Capital Flows and the Risk-taking Channel of Monetary Policy." *Journal of Monetary Economics* 71 (April): 119–32.
- Buch, C. M., M. Bussiere, L. Goldberg, and R. Hills. 2019. "The International Transmission of Monetary Policy." *Journal of International Money and Finance* 91 (March): 29–48.
- Cahill, M., S. D'Amico, C. Li, and J. Sears. 2013. "Duration Risk versus Local Supply Channel in Treasury Yields: Evidence from the Federal Reserve's Asset Purchase Announcements." Finance and Economics Discussion Series No. 35, Board of Governors of the Federal Reserve System.
- Caldara, D., and E. Herbst. 2019. "Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs." *American Economic Journal: Macroeconomics* 11 (1): 157–92.
- Campbell, J. R., C. L. Evans, J. D. Fisher, A. Justiniano, C. W. Calomiris, and M. Woodford. 2012. "Macroeconomic Effects of Federal Reserve Forward Guidance." *Brookings Papers on Economic Activity* 43 (1, Spring): 1–80.
- Cao, G., and A. Foerster. 2013. "Expectations of Large-Scale Asset Purchases." *Economic Review* (Federal Reserve Bank of Kansas City) (Q II): 5–29.
- Carlson, M. A., S. D'Amico, C. Fuentes-Albero, B. Schlusche, and P. R. Wood. 2020. "Issues in the Use of the Balance Sheet Tool." Finance and Economics Discussion Series No. 71, Board of Governors of the Federal Reserve System.
- Ca'Zorzi, M., L. Dedola, G. Georgiadis, M. Jarociński, L. Stracca, and G. Strasser. 2023. "Making Waves: Monetary Policy and

- Its Asymmetric Spillovers in a Globalized World.” *International Journal of Central Banking* 19 (2): 95–144.
- Cesa-Bianchi, A., and A. Sokol. 2022. “Financial Shocks, Credit Spreads, and the International Credit Channel.” *Journal of International Economics* 135 (March): Article 103543.
- Christiano, L. J., R. Motto, and M. Rostagno. 2014. “Risk Shocks.” *American Economic Review* 104 (1): 27–65.
- Cieslak, A., and H. Pang. 2021. “Common Shocks in Stocks and Bonds.” *Journal of Financial Economics* 142 (2): 880–904.
- Cieslak, A., and P. Povala. 2015. “Expected Returns in Treasury Bonds.” *Review of Financial Studies* 28 (10): 2859–2901.
- Cieslak, A., and A. Schrimpf. 2019. “Non-monetary News in Central Bank Communication.” *Journal of International Economics* 118 (May): 293–315.
- De Santis, R. A. 2019. “Redenomination Risk.” *Journal of Money, Credit and Banking* 51 (8): 2173–2206.
- Dedola, L., P. Karadi, and G. Lombardo. 2013. “Global Implications of National Unconventional Policies.” *Journal of Monetary Economics* 60 (1): 66–85.
- Dedola, L., and G. Lombardo. 2012. “Financial Frictions, Financial Integration and the International Propagation of Shocks.” *Economic Policy* 27 (70): 319–59.
- Dedola, L., G. Rivolta, and L. Stracca. 2017. “If the Fed Sneezes, Who Catches a Cold?” *Journal of International Economics* 108 (Supplement 1): S23–S41.
- Degasperi, R., S. S. Hong, and G. Ricco. 2024. “The Global Transmission of U.S. Monetary Policy.” Working Paper No. 1466, Bank of Italy.
- Du, W., J. Im, and J. Schreger. 2018. “The US Treasury Premium.” *Journal of International Economics* 112 (May): 167–81.
- Engel, C. 2014. “Exchange Rates and Interest Parity.” In *Handbook of International Economics*, Vol. 4, ed. G. Gopinath, E. Helpman, and K. Rogoff, 453–522. North-Holland.
- Eser, F., W. Lemke, K. Nyholm, S. Radde, and A. L. Vladu. 2023. “Tracing the Impact of the ECB’s Asset Purchase Program on the Yield Curve.” *International Journal of Central Banking* 19 (3): 359–422.
- Farhi, E., and X. Gabaix. 2016. “Rare Disasters and Exchange Rates.” *Quarterly Journal of Economics* 131 (1): 1–52.

- Fatum, R., and Y. Yamamoto. 2016. "Intra-safe Haven Currency Behavior during the Global Financial Crisis." *Journal of International Money and Finance* 66 (September): 49–64.
- 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.
- Gagnon, J., M. Raskin, J. Remache, and B. P. Sack. 2011. "Large-Scale Asset Purchases by the Federal Reserve: Did They Work?" *Economic Policy Review* 17 (1): 41.
- Galí, J., and T. Monacelli. 2005. "Monetary Policy and Exchange Rate Volatility in a Small Open Economy." *Review of Economic Studies* 72 (3): 707–34.
- Georgiadis, G., G. J. Müller, and B. Schumann. 2024. "Global Risk and the Dollar." *Journal of Monetary Economics* 144 (May): Article 103549.
- Gertler, M., and P. Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7 (1): 44–76.
- Gourinchas, P.-O., and H. Rey. 2022. "Exorbitant Privilege and Exorbitant Duty." Discussion Paper No. 16944, CEPR.
- Habib, M. M., and L. Stracca. 2012. "Getting Beyond Carry Trade: What Makes a Safe Haven Currency?" *Journal of International Economics* 87 (1): 50–64.
- Hollo, D., M. Kremer, and M. Lo Duca. 2012. "CISS — A Composite Indicator of Systemic Stress in the Financial System." Working Paper No. 1426, European Central Bank.
- Ilzetzki, E., C. M. Reinhart, and K. S. Rogoff. 2019. "Exchange Arrangements Entering the Twenty-first Century: Which Anchor Will Hold?" *Quarterly Journal of Economics* 134 (2): 599–646.
- Jarociński, M. 2022. "Central Bank Information Effects and Transatlantic Spillovers." *Journal of International Economics* 139 (November): Article 103683.
- Jarociński, M. and P. Karadi. 2020. "Deconstructing Monetary Policy Surprises—The Role of Information Shocks." *American Economic Journal: Macroeconomics* 12 (2): 1–43.
- Jayawickrema, V., and E. Swanson. 2021. "Speeches by the Fed Chair Are More Important than FOMC Announcements: An Improved High-frequency Measure of US Monetary Policy Shocks." Unpublished Manuscript.

- Jiang, Z., A. Krishnamurthy, and H. Lustig. 2021. "Foreign Safe Asset Demand and the Dollar Exchange Rate." *Journal of Finance* 76 (3): 1049–89.
- Kearns, J., A. Schrimpf, and F. Xia. 2023. "Explaining Monetary Spillovers: The Matrix Reloaded." *Journal of Money, Credit and Banking* 55 (6): 1535–68.
- Krishnamurthy, A., and A. Vissing-Jorgensen. 2012. "The Aggregate Demand for Treasury Debt." *Journal of Political Economy* 120 (2): 233–67.
- Maggiore, M., B. Neiman, and J. Schreger. 2019. "The Rise of the Dollar and Fall of the Euro as International Currencies." *AEA Papers and Proceedings* 109 (May): 521–26.
- Matheson, T., and E. Stavrev. 2014. "News and Monetary Shocks at a High Frequency: A Simple Approach." *Economics Letters* 125 (2): 282–86.
- 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.
- Miranda-Agrippino, S., and T. Nenova. 2022. "A Tale of Two Global Monetary Policies." *Journal of International Economics* 136 (May): Article 103606.
- Miranda-Agrippino, S., and H. Rey. 2020. "US Monetary Policy and the Global Financial Cycle." *Review of Economic Studies* 87 (6): 2754–76.
- Nakamura, E., and J. Steinsson. 2018. "High-frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics* 133 (3): 1283–1330.
- Pastor, L., and P. Veronesi. 2009. "Learning in Financial Markets." *Annual Review of Financial Economics* 1: 361–81.
- Pinchetti, M., and A. Szczepaniak. 2024. "Global Spillovers of the Fed Information Effect." *IMF Economic Review* 72 (2): 773–819.
- Rigobon, R., and B. Sack. 2003. "Measuring the Reaction of Monetary Policy to the Stock Market." *Quarterly Journal of Economics* 118 (2): 639–69.
- . 2004. "The Impact of Monetary Policy on Asset Prices." *Journal of Monetary Economics* 51 (8): 1553–75.
- Rostagno, M., C. Altavilla, G. Carboni, W. Lemke, R. Motto, A. Saint Guilhem, and J. Yiangou. 2021. *Monetary Policy in Times*

of Crisis: A Tale of Two Decades of the European Central Bank.
Oxford University Press.

- 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.
- Stock, J. H., and M. W. Watson. 2012. "Disentangling the Channels of the 2007-2009 Recession." *Brookings Papers on Economic Activity* 43 (1, Spring): 81–156.
- Vissing-Jorgensen, A., and A. Krishnamurthy. 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." *Brookings Papers on Economic Activity* 42 (2, Fall): 215–87.

Monetary Policy Strategies to Navigate Post-Pandemic Inflation: An Assessment Using the ECB's New Area-Wide Model*

Matthieu Darracq Pariès, Antoine Kornprobst,
and Romanos Priftis
European Central Bank

We assess the effectiveness of ECB monetary policy in the post-pandemic period using a state-of-the-art estimated DSGE model of the euro area. Through counterfactual experiments varying the timing of interest rate hikes, policymaker preferences, and optimality criteria, we compare alternative strategies against actual policy. The ECB's decision to begin hiking the policy rate in 2022:Q3 effectively balanced inflation and output stabilization, leading to a brief slowdown and inflation converging to target by end-2025. An earlier hike in 2021:Q4 would have reduced peak inflation but caused a prolonged recession, while a delayed hike to 2024:Q1 would have fueled persistent inflation above 4 percent until 2026. Optimal policy simulations indicate that a more aggressive tightening could have lowered inflation volatility only at the cost of deeper contractions in GDP. Real-time monetary policy counterfactuals confirm that actual policy closely aligned with optimal prescriptions, effectively managing the trade-off between inflation and growth.

JEL Codes: C53, E31, E42, E52, E58.

*The views expressed in this paper are those of the authors and should not be attributed to the European Central Bank or the Eurosystem. We thank the Editor and two anonymous referees for valuable suggestions and comments as well as seminar and conference participants at the European Central Bank. Author emails: matthieu.darracq-paries@ecb.europa.eu, antoine.kornprobst@ecb.europa.eu, romanos.priftis@ecb.europa.eu.

1. Introduction

In response to increasing inflation in 2021, the European Central Bank (ECB) took key actions to unwind its accommodative monetary policy stance. In terms of its primary instruments, 2022:Q3 marked the end of an extended period where the deposit facility rate stood near the effective lower bound since the outbreak of the sovereign debt crisis in 2012. The short-term interest rate was raised from -0.5 percent in 2022:Q3 and reached 4.0 percent four quarters later in 2023:Q3. Notwithstanding the tightening in monetary policy, inflation in the euro area peaked at 10 percent in 2022:Q4 and has been declining since then, to 2.7 percent in 2023:Q4. According to the projections published by the ECB in December 2023, this figure is expected to converge to the 2 percent medium-term target in 2025.

How successful was ECB monetary policy at stabilizing euro-area inflation and GDP in the post-pandemic period? We employ the workhorse dynamic stochastic general equilibrium (DSGE) model used for policy analysis at the ECB—the New Area-Wide Model II (NAWM II) (Coenen et al. 2018)—to assess the effectiveness of the ECB’s monetary policy in stabilizing inflation and real activity during this volatile post-pandemic period. In particular, we conduct alternative monetary policy experiments, which differ in their timing of interest rate tightening, the specification of policymaker preferences, and their optimality with respect to stabilization and welfare-based metrics. Our intention is to contrast the effectiveness of these alternative monetary policy strategies against the impact of actual policy against using a state-of-the-art DSGE policy model tailored to provide quantitative prescriptions regarding euro-area monetary policy.

Our results suggest that the ECB was relatively effective in stabilizing inflation and output, though not without trade-offs. The model-based analysis suggests that the actual policy path—whereby the hiking started in 2022:Q3—resulted in a brief economic slowdown, with inflation converging to target by end-2025, and a peak rate of 4 percent. An earlier rate liftoff in 2021:Q4 would have reduced peak inflation (8 percent versus 10 percent) but at the cost of a prolonged recession. A delayed liftoff to 2024:Q1 would have boosted growth but led to more persistent inflation (above

4 percent until 2026). Optimal policy simulations indicate that a more hawkish stance could have reduced inflation volatility but at the expense of a deeper recession. Real-time counterfactuals confirm that actual policy closely followed optimal prescriptions, striking an efficient balance between inflation and output stabilization. Most alternative strategies, except for a delayed tightening, improved price stability but increased output volatility, reinforcing the notion that the ECB effectively managed the trade-off between inflation and growth.

Our setup builds on recent studies that assess counterfactual monetary policy strategies during the post-pandemic inflation period, extending the analysis from the U.S. to the euro area. For instance, Crump et al. (2023) use the Federal Reserve Bank of New York's DSGE policy model to simulate counterfactual monetary policy strategies of the Federal Reserve, while De Fiore et al. (2023) compare actual U.S. monetary policy to alternative monetary policy frameworks using an estimated policy model. Our work extends these studies by offering a model-based evaluation of key ECB monetary policy decisions of the recent years using the ECB's estimated policy model. In particular and differently to others in the literature, it allows for policy to be counterfactually set in an optimal manner and to evaluate its effectiveness by changing the information set of the central bank, that is by taking a perfect-foresight perspective where all future realized shocks are known with certainty, or a real-time approach where only past shocks are observed by the central bank and future shocks are derived from vintages of official staff projections. In this regard, our work is related to several other studies analyzing the effects of monetary policy rules in real time, such as Orphanides (2001), Molodtsova, Nikolsko-Rzhevskyy, and Papell (2008), and Altavilla and Ciccarelli (2011), among others, although, to the best of our knowledge, this is the first paper to assess the real-time effects of *optimal* monetary policy effectiveness using a large-scale structural model.

In a first step, we ask the counterfactual question of whether the timing of the hike mattered. In particular, with hindsight of the evolving macroeconomic conditions and conditional on the most recent macroeconomic projections, how effective was euro-area monetary policy in stabilizing inflation by choosing 2022:Q3 to raise interest rates? And what would have happened to interest rates,

output, and inflation had the ECB chosen to begin its tightening cycle at different points in time since 2021:Q4 when inflationary pressures started to materialize?

According to the NAWM II, aligning the timing of the liftoff with actual policy decisions—starting in 2022:Q3—suggests that interest rate hikes would have initially been somewhat steeper, but tightening would have proceeded more gradually thereafter, with rates peaking at 3.75 percent in 2024:Q2 rather than 4 percent as observed in practice. In this event, the economic slowdown would have been brief, but inflation in 2024 would have been higher by 0.3 percentage points and converged to the 2 percent target by the end of 2025. Instead, an earlier interest rate hike beginning in 2021:Q4 would have limited the inflation peak to 8 percent rather than 10 percent albeit at the cost of a protracted recession in 2022 and 2023. By contrast, a counterfactual where monetary tightening is delayed to 2024:Q1—representing a strong commitment to “look through” the sequence of inflationary supply shocks—would have resulted in significantly higher economic growth but a more entrenched inflationary environment, where inflation would have peaked around 10.5 percent and remained above 4 percent until 2026.

While the first set of counterfactual strategies assumes monetary policy follows the estimated interest rate rule of NAWM II, a natural question arises: what would the effects have been if, for each hike’s starting date, monetary policy had been designed optimally? To answer this question we compute a set of optimal interest rate paths at different points in time, determined by means of minimizing an intertemporal loss function conditional on the latest baseline projections available for the euro area. Our application is similar in spirit to other studies employing large-scale DSGE models to assess optimal monetary policies under commitment using predetermined loss function specifications; for example, see Adjemian, Darracq Pariès, and Moyen (2007); Adolfson et al. (2011); Darracq Pariès and Kühn (2016); de Groot et al. (2021); Harrison and Waldron (2021); and Mazelis, Motto, and Ristinieniemi (2023). In particular, our analysis is complementary to parallel work by Carrier and Mavromatis (2024), who analyze optimal normalization strategies during the post-pandemic period in a setting where inflation expectations may become de-anchored. While their behavioral model explicitly captures nonlinearities arising from expectation dynamics, our model

prioritizes empirical fit by allowing for policy-relevant transmission channels and estimating its parameters based on observed data.

To solve for the optimal policy in the NAWM II, we employ the constrained optimal policy projection from de Groot et al. (2021). Harrison and Waldron (2021), Barnichon and Mesters (2023), McKay and Wolf (2023), Carrier and Mavromatis (2024), and Dengler et al. (2024) have employed similar frameworks in related studies. Compared to the rule-based counterfactuals, the optimal policy counterfactuals are more hump-shaped, with higher peak rates but less persistence of the restrictive stance. Here, picking 2022:Q3 as the initial date of monetary tightening would engineer a mild recession and deliver inflation outcomes similar to those observed in the data and under the rule-based counterfactual. The more hawkish preferences would imply a deeper recession and inflation undershooting the target by 2024 already.

Although the retrospective analysis yields important conclusions about the effectiveness of euro-area monetary policy, in practice, policy decisions are made in real time, based on the information set available at each point in time, rather than with the benefit of hindsight. As a result, policy mistakes can arise not only from misjudging the economic outlook but also from failing to react promptly to new developments. Our third set of exercises therefore relaxes the assumption of perfect information from the perspective of the central bank, and assesses the effectiveness of an optimal *real-time* monetary policy. By analyzing real-time policy counterfactuals, we assess how policymakers responded to shocks as they unfolded, distinguishing between the impact of unforeseen events and the role of policy choices in shaping macroeconomic outcomes. This approach—which to our knowledge is novel to the literature from the perspective of large-scale structural models—provides a more realistic evaluation of monetary policy effectiveness than assuming perfect foresight of the baseline projection and shocks.

In this setting, we update optimal monetary policy counterfactuals each quarter between 2021:Q4 and 2023:Q4, reflecting new vintages of incoming data and macroeconomic projections. Our results suggest that, unlike in simulations with perfect foresight about the sequence of shocks hitting the economy, in real time there is no strong rationale for hiking policy rates aggressively already in 2022:Q1. Instead, real-time prescriptions call for a less steep

interest rate increase, which is however more persistent than in full-information counterfactual simulations. The real-time simulations confirm that actual policy fares relatively well against alternative policy conducts. In particular, the actual rate path aligns more closely with real-time optimal policy projections than with the model's estimated interest rate rule.

Given the numerous monetary policy strategies evaluated so far, we compare stabilization and welfare metrics implied by each strategy over time. We find that most alternative strategies, with the notable exception of an extended lower-for-longer policy where the hiking cycle is initiated in 2024:Q1, improve price stability relative to actual policy. That said, no strategy improves both GDP and inflation volatility relative to actual policy, implying that inflation stabilization can only be improved at the cost of increased output volatility. In terms of a normative comparison, optimal policy strategies are generally associated with lower welfare losses relative to actual policy, in particular under more hawkish preferences of the central bank and when initiating the hike earlier. Overall, the proximity of the benchmark optimal policy macroeconomic allocation with the observed macroeconomic baseline indicates that the actual policy conduct managed to implement an efficient management of the output-inflation trade-off.

It is important to recognize that the validity of these simulation results, and their corresponding welfare assessment, depends on the model's underlying assumptions. Specifically, the analysis relies on a first-order approximation to describe the model's dynamics. Given the large shocks observed during the periods examined, some nonlinear effects may not be fully captured, potentially leading to less persistent impacts or muted responses, thereby understating the welfare losses associated with different strategies. Furthermore, the model assumes that longer-term inflation expectations remained anchored throughout the analysis period. While this assumption appears reasonable in hindsight, it may influence the results by dampening potential feedback effects on inflation dynamics and, consequently, welfare, particularly in times of heightened uncertainty.

We take steps to mitigate uncertainty around key model assumptions by conducting robustness checks on the strength of policy transmission—specifically, through variations in the slope of price and wage Phillips curves and the degree of financial frictions. These

sensitivity analyses help address concerns about nonlinearity by exploring different regimes in which policy transmission operates differently. We find that under a steeper Phillips curve, a lower increase in the policy rate would be sufficient to limit the peak of inflation relative to the benchmark case, while output would have decreased by less. Moreover, under a stronger degree of financial frictions (captured by an increase in households' portfolio adjustment costs) the central bank faces a larger sacrifice ratio, leading similar interest rate hikes to produce disinflation but at a higher output cost. Finally, we also highlight the importance of agent expectations in shaping the predictions of optimal policy. We find that expectations influence the optimal interest rate trajectory with policy announcements further away in the future being less effective when deviating from rational expectations.

The remainder of this paper is structured as follows. Section 2 introduces the ECB's estimated DSGE model and provides an overview of its key features. Section 3 illustrates the several counterfactual monetary policy strategies, while Section 4 performs robustness checks around key transmission mechanisms. An appendix provides additional figures and information on extensions.

2. An Overview of the Model

We use as a laboratory the NAWM II, the workhorse estimated DSGE policy model of the ECB, which has been designed for use in the Eurosystem's (Broad) Macroeconomic Projection Exercises (BMPE) and for analysis of topical policy issues affecting business cycle fluctuations.¹ The model comprises the euro area and a rest-of-the-world block. It is estimated using Bayesian techniques and the list of observables includes key financial variables, in particular bond yields and bank lending rates. In light of the large number of time series used for estimation, the model features 24 distinct structural shocks. In what follows we provide an overview of the model's

¹See, e.g., Cecion et al. (2021) and Brand et al. (2021) (and references therein) for recent applications of the NAWM II model related to price stability, the inflation-output trade-off, among other key topical issues. See also Ciccarelli, Darracq Pariès, and Priftis (2024) for an overview of ECB policy uses and applications using the NAWM II.

main features, with a focus on outlining key relationships within the financial sector.²

NAWM II is an open-economy DSGE model with financial intermediaries and features seven types of economic agents: households, firms, wholesale banks, retail banks, capital producers, a fiscal authority, and the central bank.

Households. Households make optimal choices regarding their purchases of consumption and investment goods, the latter determining the economy-wide capital stock. They supply differentiated labor services in monopolistically competitive markets, and set wages as a mark-up over the marginal rate of substitution between consumption and leisure.

In terms of financial assets, households can place deposits in the amount of $D_{h,t+1}$ to wholesale banks and receive a (gross) return $R_{D,t-1}$ on previous-period deposits. They can also hold short-term (one-period) government bonds, B_{t+1} , and invest in domestic long-term government bonds subject to portfolio adjustment frictions:

$$\Gamma_L^h(B_{L,h,t+1}) = 1/2 \gamma_L^h (B_{L,h,t+1} - \bar{B}_{L,h})^2 / \bar{B}_{L,h}. \quad (1)$$

For purchasing investment goods, $I_{h,t}$, households demand loans, which are originated by wholesale banks and distributed by retail banks. At the retail level, loans are nominal consoles with geometrically decaying coupons (as in Woodford 2001), with $\mathcal{Q}_{I,t} B_{I,t+1}$ denoting the loan value and $\mathcal{R}_{I,t} = \frac{1}{\mathcal{Q}_{I,t}} + \varrho_I$ denoting the loan rate (“yield to maturity”), where $\mathcal{Q}_{I,t}$ is the discount price of the loan $B_{I,t+1}$ and ϱ_I refers to the decay factor of the coupons. Demand for loans arises as households face a “loan-in-advance” (LIA) financing constraint for investments in physical capital (as in Carlstrom, Fuerst, and Paustian 2017):

$$P_{C,t} \tilde{p}_{I,t} I_{h,t} \leq \mathcal{Q}_{I,t} (B_{I,t+1} - \varrho_I B_{I,t}), \quad (2)$$

where $\tilde{p}_t^I$ is the relative price of the investment good obtained from the capital goods producer. In turn, the value of the loan

²For a detailed description of the model’s structure, see Coenen et al. (2018), while for a detailed exposition of the ECB macroeconomic projection exercise, see ECB (2016).

for household h can be expressed as $\tilde{p}_{I,t} (1 + \varsigma_{h,t}) = Q_{h,t} \epsilon_t^I$, where $\varsigma_{h,t} \Lambda_{h,t}/P_{C,t}$ is the Lagrange multiplier associated with the household's LIA constraint and $Q_{h,t}$ represents Tobin's Q .

Wholesale and Retail Banks. The investment loans are supplied by wholesale and retail banks, which act as financial intermediaries. First, wholesale banks use the deposits raised from the households and their net worth to issue wholesale loans and to purchase domestic and foreign long-term bonds. Second, retail banks purchase the wholesale loans, transform them into retail loans, and offer these retail loans to the households. Because of the financing constraint on the part of the households, the production of new physical capital is separated from the households' decision problem and relegated to capital producers.

In particular, banks' assets comprise long-term loans to the private sector for investment purposes ($B_{I,b,t+1}^p$), domestic ($B_{L,b,t+1}^p$) and foreign ($B_{L,b,t+1}^{*,p}$) government bonds, and reserves held at the central bank. On the liabilities side, they finance themselves via the collection of short-term deposits and retained earnings. Their balance sheet is hence given by

$$\tilde{Q}_{I,t} B_{I,b,t+1}^p + Q_{L,t} B_{L,b,t+1}^p + S_t Q_{L,t}^* B_{L,b,t+1}^{*,p} = NW_{b,t} + D_{b,t+1}^h. \quad (3)$$

Reflecting an agency problem where banks must hold sufficient capital to reassure their creditors that they will not divert funds, banks are subject to a (risk-weighted) leverage constraint (see Gertler and Karadi 2011, 2013). Their leverage, Φ_t , is therefore constrained by their net worth, NW_t :

$$\tilde{Q}_{I,t} B_{I,t+1}^p + \omega_L Q_{L,t} B_{L,t+1}^p + \omega_{L,t}^* S_t Q_{L,t}^* B_{L,t+1}^{*,p} = \Phi_t NW_t,$$

where $\tilde{Q}_{I,t}$, $Q_{L,t}$, and $Q_{L,t}^*$ are the "discount prices" of the privately (p) intermediated assets $B_{I,t+1}^p$, $B_{L,t+1}^p$, and $B_{L,t+1}^{*,p}$. The capital position of banks thus influences the response of the economy to shocks via a financial accelerator mechanism. Incorporating adjustment costs into the pricing of loans, the model also includes the sluggish adjustment of lending rates.

Firms. As regards firms, the model distinguishes between domestic producers of tradable intermediate goods and domestic

producers of three types of nontradable final goods: a private consumption good, a private investment good, and a public consumption good. The intermediate goods firms use labor and capital services as inputs to produce differentiated goods, which are sold in monopolistically competitive markets domestically and abroad. In doing so, they set different prices for domestic and foreign markets as a markup over their marginal costs. The final goods firms combine domestic and foreign intermediate goods in different proportions, acting as price takers in fully competitive markets. The foreign intermediate goods are imported from producers abroad, who set their prices in euros in monopolistically competitive markets, allowing for a gradual exchange rate pass-through. A foreign retail firm in turn combines the exported domestic intermediate goods, with aggregate export demand depending on total foreign demand.

Besides the financial frictions described above, both households and firms face a number of nominal and real frictions, which have been identified as important in generating empirically plausible dynamics. Additional real frictions are introduced via external habit formation in consumption, through generalized adjustment costs in investment, imports and exports, and through fixed costs in intermediate goods production. Nominal frictions arise from staggered price and wage setting as in Calvo (1983), in combination with (partial) dynamic indexation of price and wage contracts to past inflation.

Government. The fiscal authority purchases the public consumption good, issues domestic bonds, and levies different types of distortionary taxes, albeit at constant rates; Ricardian equivalence holds because the fiscal authority's budget is balanced each period using lump-sum taxes.

Monetary Policy. Given our focus on monetary policy strategies, we describe the model interest rate rule, which we make use of in our initial counterfactual exercises of Section 3. Specifically, the monetary authority is assumed to follow an inertial Taylor-type interest rate rule targeting the response of inflation, the output gap, and their respective quarterly changes:

$$i_t = \rho i_{t-1} + (1 - \rho) (\bar{r} + \bar{\pi} + \phi_{\hat{\pi}} \hat{\pi}_t + \phi_{\Delta\pi} \Delta\pi_t + \phi_{\hat{y}} \hat{y}_t + \phi_{\Delta\hat{y}} \Delta\hat{y}_t) + \eta_t. \quad (4)$$

Table 1. Taylor Rule Estimated Parameters

Variables		Posterior Mode
ρ	Interest Rate Smoothing	0.9346
$\phi_{\hat{\pi}}$	Response to Inflation	2.9334
$\phi_{\Delta\pi}$	Response to change in Inflation	0.0361
$\phi_{\hat{y}}$	Response to Output Gap	0.0320
$\phi_{\Delta y}$	Response to Change in Output Gap	0.0920

The parameters of the interest rate rule (see Table 1) are estimated on euro-area data over the 1985:Q1–2019:Q4 period, using Bayesian methods.

Finally, the model is closed by a rest-of-the-world block, which is represented by a structural vector autoregressive (SVAR) model determining five foreign variables: foreign demand, foreign prices, the foreign interest rate, foreign competitors’ export prices, and the price of oil.

3. Counterfactual Monetary Policy Strategies

We assess how the euro-area economy would have performed had monetary policy deviated from actual interest rate policy and considered a different starting period of its hiking cycle by following the model-consistent interest rate rule. We contrast these results against the predictions from an interest rate policy, which is designed optimally. We conduct both *ex post* and *real-time* assessments of various monetary policy approaches by changing the information set of the central bank in relation to its knowledge about future incoming data.

3.1 Macroeconomic Projection Baseline

A central element of our counterfactual experiments is the ECB/Eurosystem’s staff projection baseline of euro-area macroeconomic variables, and the extent to which agents in our model are aware of it. As part of the Eurosystem projection process (i.e., the Broad Macroeconomic Projection Exercises, or BMPE), such a baseline comprises quarterly paths for euro-area real GDP growth, HICP

(Harmonised Index of Consumer Prices) inflation, and other macroeconomic aggregates. In particular, the official Eurosystem projection baseline is constructed using exogenous assumptions of the international economic environment, as well as an aggregation of national-level fiscal policies. Interest rates are assumed to follow the euro short-term rate (€STR) forward curve, making the baseline consistent with market expectations of monetary policy conduct. These assumptions are combined with additional off-model information, including expert judgment about the transmission and persistence of shocks.³

We treat such a projection baseline as an input into the model, for the purpose of rationalizing and projecting macroeconomic aggregates over the simulating horizon. Specifically, throughout the experiments we perform a series of model-based forecasts which condition on all structural shocks filtered from each baseline, other than the standard monetary policy shock.

Notably, our approach consists of evaluating the model-implied interest rate policy (either based on the estimated monetary policy rule or in an optimal manner), taking other policies and shocks as given. In particular, the structural shocks filtered from the baseline also contain information on the fiscal policy response, which is fully captured in the shock decomposition. As such, we abstract from the jointly optimal fiscal–monetary policy interaction. Moreover, the baseline for macroeconomic variables also incorporates the effects from other parallel monetary policies involving other instruments that were also implemented from 2022 onward, such as the termination of net asset purchases in 2022:Q3, the introduction of the Transmission Protection Instrument (TPI) to support the effective transmission of monetary policy across countries, and the recalibration of targeted longer-term refinancing operations (TLTRO III).

3.2 Timing of the Hike: What if the ECB Followed the Estimated Taylor Rule?

We place ourselves in 2021:Q4, 2022:Q3, or 2024:Q1 and assume full knowledge of incoming data and the Eurosystem’s staff off-model

³ More detailed illustrations of the use of NAWM II for forecasting and policy analysis can be found in Ciccarelli, Darracq Pariès, and Priftis (2024).

baseline forecast from the December 2023 projection.⁴ We then determine the interest rate path, as well as the associated macro-economic allocation, that would be consistent with the historical monetary policy reaction function as implied by our model.


The rule-based projection is implemented by performing a conditional forecast and conditioning on all structural shocks filtered from the December 2023 baseline, other than the standard monetary policy shock. This allows taking into account all non-monetary structural shocks that, according to the model, had hit or were expected to hit the economy between the starting date of the counterfactual simulation and 2026:Q4, including global supply shocks, energy and food price shocks, and the tightening of the labor market. To the extent that the estimated rule of the model is a good approximation of market expectations about ECB interest rate policy, deviations of historical and forward interest rates from the rule-based projection would reflect the discretionary component of monetary policy interventions.

Figure 1 depicts the rule-based policy counterfactuals whereby the ECB ended its lower-for-longer policy and followed the historical reaction function since 2021:Q4 (solid light blue lines) and since 2022:Q3 (dashed light blue lines), or extended the lower-for-longer policy until 2023:Q4 (dotted light blue lines), compared with historical data and baseline projections (solid dark blue lines).

Had the ECB followed the rule-based policy in 2021:Q4, it would have hiked the short-term interest rate immediately, reaching a level of 3.75 percent in 2022:Q4, which it would maintain roughly until end-2025. Compared to actual policy, the counterfactual tightening impulse would have initially slowed GDP growth more markedly and the euro-area economy would have been in a recession in 2022 and the first half of 2023. Annual HICP inflation, in turn, would have been lower, peaking around 8 percent instead of 10 percent in the end of 2022 due to the sharp slowdown in economic activity. Relative to the baseline, inflation would have reached the 2 percent target eight quarters earlier, in 2024:Q1 and been projected to settle below the target in 2025.

⁴In Section 3.4 we relax the assumption of full information and repeat the exercise in real time, whereby in each period only the contemporaneous vintage of each baseline is known.

Figure 1. Rule-Based Counterfactuals in 2021:Q4, 2022:Q3, and 2024:Q1



Note: This figure depicts rule-based counterfactual policies whereby the ECB follows the NAWM policy rule since 2021:Q4 (solid light blue lines), 2022:Q3 (dashed light blue lines), or 2024:Q1 (dotted light blue lines), conditioning on the data and macroeconomic outlook from Eurosystem’s staff projections in December 2023 (solid dark blue lines).

Instead, if the lower-for-longer period ended in 2022:Q3 (dashed light blue lines)—in line with actual policy conduct—and the model-consistent rule-based policy was followed thereafter, the interest rate hikes would have been somewhat higher than in the actual policy path until 2023:Q1, but with more gradual tightening thereafter, only reaching a peak of 3.75 percent in 2024:Q2. The profile of this rule-based policy projection indicates that the actual hiking cycle can be divided into two phases: first, a swift normalization of policy rates to prevent a protracted decline in real interest rates from boosting demand and fueling additional inflationary pressures; second, gradual adjustments towards the peak rate and a restrictive stance to further anchor expectations at the 2 percent target. The effects under this policy counterfactual suggest that the economic slowdown would have been short lived—after a quarter below zero, GDP growth would have rebounded more markedly in 2023:Q2—with a similar disinflationary process as in the baseline for 2023 but higher inflation in 2024 by 0.3 percentage points. The convergence to the 2 percent target would be reached by end-2025 as in the baseline.

Last, we contrast the effects of these policies against the counterfactual whereby the lower-for-longer policy is maintained until 2023:Q4 (dotted light blue lines) before following the model-consistent rule-based policy from 2024:Q1 onward. The counterfactual can be interpreted as the strongest form of “looking through” the full sequence of supply shocks, leaving policy rates unchanged until the beginning of 2024. In this counterfactual, the lack of monetary tightening would have pushed output significantly above potential between 2022:Q2 and 2025:Q1 and led to higher inflation expectations, more pronounced second-round effects, and thus higher and more persistent price inflation, peaking around 10.5 percent and remaining above 4 percent until 2026.

Overall, the results suggest that with full information about the macroeconomic conditions that would prevail since 2021, an earlier hike would have reduced the peak of inflation and achieved convergence to the target faster, but would have driven the euro-area economy into recession all through 2022. However, the 2022:Q3 counterfactual exercise demonstrates that, once aligning the interest rate liftoff date with actual policy conduct, the counterfactual monetary tightening would be broadly similar to the baseline in terms of inflation and interest paths. From this standpoint, the estimated interest

rate rule of the NAWM II tracks relatively well the tightening cycle observed from 2022:Q3 to 2024:Q1. Keeping interest rates unchanged for longer, as in the 2024:Q1 counterfactual, would lead to higher inflation rates becoming entrenched with significant second-round effects materializing over the medium term.

3.3 What Would Optimal Interest Rate Policy Prescribe?

While conditional forecasts using the model-consistent interest rate rule enable assessing the effects of different starting points for the hiking cycle, a natural question that arises is whether these counterfactual monetary policy strategies are optimal from a loss function minimization perspective.

To address this, we compute optimal policy counterfactuals—at various starting dates—under commitment and conditional on the December 2023 macroeconomic baseline and a standard central bank loss function. We assume that the central bank commits ex ante to setting the short-term interest rate to minimize a quadratic loss function that approximates household utility, penalizing deviations of inflation from target, output from potential, and excessive interest rate volatility.

The approach builds on the concept of optimal policy projections under commitment, first introduced by Svensson (2005), who formalized the selection of alternative paths for target variables and policy instruments to minimize a central bank's loss function. It was first applied in practice using the Federal Reserve's FRB/US model by Svensson and Tetlow (2005), and is closely related to the targeting rule framework for optimal policy by, e.g., Svensson and Woodford (2003). Specifically, among the set of alternative projections that shift the target and policy variables away from baseline paths using anticipated and unanticipated impulse responses to the policy instrument, the method relies on identifying the projection that minimizes the policymaker's objective function. In our implementation, we use the COPPs (constrained optimal policy projections) toolkit of de Groot et al. (2021) in Dynare (Adjemian et al. 2022). Notably, complementary methods have been proposed by Barnichon and Mesters (2023), who show that the gradient of the central bank's loss function in response to policy shocks is a sufficient statistic for capturing the optimal policy, and also by McKay and Wolf (2023), who

leverage sequence-space representations of equilibria to show that in linearized structural macroeconomic models, the impulse responses to policy news shocks are sufficient to construct policy rule counterfactuals. Carrier and Mavromatis (2024) extend the methods to a behavioral model whereby expectations may become de-anchored.

Our assumed loss function follows the standard specification found in the optimal monetary policy literature (see, e.g., Giannoni and Woodford 2003, Woodford 2010, Debortoli et al. 2019, or Harrison and Waldron 2021). Formally, the loss for strategy j , between t and $t + h$ is defined by

$$\mathcal{L}_j = \frac{1}{2} \sum_t^{t+h} \beta^t \mathbb{E}_t (\lambda_\pi \hat{\pi}_{t,j}^2 + \lambda_y \hat{y}_{t,j}^2 + \lambda_{\Delta i} \Delta i_{t,j}^2), \quad (5)$$

where β is the discount factor, $\hat{\pi}$ is the difference between inflation and the inflation target, $\hat{y}$ is the percent deviation of output from its potential, and Δi is the quarterly change in the short-term nominal interest rate.

The relative weights attached to inflation, output, and rate stabilization can be interpreted as reflecting central bank policy preferences. Being agnostic on these values, we calibrate them such that the “benchmark” optimal policy projection starting in 2022:Q3 implies a short-term rate whose path approximates the baseline path of the euro short-term rate (€STR) forward curve. A more “hawkish” loss function would then require a lower weight on output stabilization.⁵ Starting from a value of $\lambda_\pi = 1$, the benchmark calibration yields $\lambda_y = 0.5$ and $\lambda_{\Delta i} = 4$. Our “hawkish” calibration considers a stronger focus on price stability with a lower weight on output $\lambda_y = 0.1$.⁶

⁵As with the rule-based policy counterfactual, we assume that the policy-maker adjusts only the short-term interest rate, while central bank balance sheet runoffs proceed in line with market expectations and their effects are included in the macroeconomic projection baseline. Similarly, we abstract from the jointly optimal determination of monetary–fiscal policy and take fiscal policy as given, with its effects being incorporated into shaping the macroeconomic baseline.

⁶Note that in Section 4.1, we perform more general sensitivity analysis around the relative weights attached to the output gap, λ_y , and interest rate smoothing, $\lambda_{\Delta i}$, which allows computing a range of optimal policy projections under the constellation of policy preferences consistent with medium-term price stability.


Figure 2 depicts optimal policy counterfactuals at different points in time and under both the “benchmark” and “hawkish” central bank preferences. Compared to the estimated interest rate rule-based counterfactual, optimal policy counterfactuals are more hump-shaped, with higher peak rates but less persistence of the restrictive stance. The benchmark preferences would engineer a mild recession and would deliver inflation outcomes by 2024 and 2025 quite similar to the baseline. Instead, more hawkish preferences would engineer a deeper recession and be associated with a faster disinflationary process followed by a mild undershooting of the target in 2024.⁷

Specifically, the optimal policies announced in 2021:Q4 (solid red or yellow lines) prescribe hiking the policy rate gradually from -0.5 percent in 2021:Q4 to between 3.5 percent and 4.5 percent, respectively, in the period 2023:Q2–2023:Q4. The peak rate is lower and is reached later in 2023 if the output weight in the loss function is high (solid yellow lines), whereas it is higher and reached earlier in 2023 if the output weight is smaller (solid red lines). At the same time, optimal policy with hawkish preferences generates a significant disinflation and implies rate cuts already in mid-2023 to avoid a sharp undershooting of inflation. Optimal policy with benchmark preferences, in turn, would necessitate a lower but more persistent peak rate. In the two optimal policies, the early start of the hiking cycle achieves to limit the peak of inflation but at the expense of more severe downturn in GDP growth.

In turn, when optimal policy is implemented in 2022:Q3 (dotted red and yellow lines), the paths for the policy rate are less gradual and their peaks range between 4 percent and 4.25 percent, reflecting that the higher inflation rates have already materialized at the time of policy shift. By calibration design, the benchmark optimal policy implies a similar interest rate path as in the baseline but achieves marginally lower inflation rates all through the simulation horizon, at the cost of lower GDP growth rates for few quarters in the turn of

⁷Given the hump-shaped nature of the optimal interest rate paths, Appendix A.1 extends the analysis and compares the optimal policy outcome under alternative average inflation targeting (AIT) frameworks. Relatedly, in Section 3.5 we illustrate how the optimal policy can be implemented by optimizing on the coefficients of the interest rate rule of the NAWM II, following the approach in Schmitt-Grohe and Uribe (2007).

Figure 2. Optimal Policy Projections in 2021:Q4 and 2022:Q3



Note: This figure depicts counterfactual optimal policy projections at different points in time, under alternative parameterizations of the central bank loss function and conditioning on the data and macroeconomic outlook from Eurosystem’s staff projections in December 2023. Solid red (yellow) lines represent the optimal policy in 2021:Q4 with hawkish (benchmark) preferences. Dotted red (yellow) lines represent the optimal policy in 2022:Q3 with hawkish (benchmark) preferences.

2023 (dotted yellow line). The proximity of the benchmark optimal policy macroeconomic outcomes with the baseline, once assuming the 2022:Q3 liftoff date, brings further indication that the actual policy conduct managed to implement an efficient management of the output-inflation trade-off. By comparison, the rate hikes under the hawkish optimal policy, given the later start of the hiking cycle, would peak above 4 percent and remain persistently in restrictive territory (dotted red line). This would be sufficient to bring inflation back to target by 2024:Q1 but at the expense of a more pronounced and persistent reduction in GDP growth relative to baseline.

3.4 Monetary Policy in Real Time: Bad Luck or Bad Policy?


The medium-term orientation of monetary policy implies a forward-looking approach that relies heavily on the economic outlook, making accurate macroeconomic projections essential for effective policy calibration. In reality, however, unforeseen macroeconomic shocks often lead to forecast errors and, consequently, suboptimal policy decisions when judged *ex post*. A prominent example is the Russian invasion of Ukraine in 2022:Q1, which triggered lasting supply disruptions and a surge in energy prices. Had the ECB foreseen this event and its economic consequences back in December 2021, an earlier interest rate hike may have been warranted.

We now simulate policy counterfactuals in real time to assess the respective role of “bad luck” (i.e., one-sided repeated adverse shock hitting the euro-area economy) and “bad policy” (i.e., delayed monetary tightening or suboptimal management of the output-inflation trade-off) in shaping macroeconomic outcomes since 2022:Q1. In this setting, policy counterfactuals are updated each quarter between 2021:Q4 and 2023:Q4, reflecting new vintages of incoming data and macroeconomic projections. This exercise illustrates how policy projections are revised quarter after quarter as new shocks hit the economy over that period.⁸

Figure 3 illustrates the comparison between actual policy and a real-time staggered policy setting, using either the interest rate

⁸The importance of real-time data for estimating monetary policy rules and evaluating the effectiveness of monetary policy is also described in Orphanides (2001), Molodtsova, Nikolsko-Rzhevskyy, and Papell (2008), and Altavilla and Ciccarelli (2011), among others.

Figure 3. Staggered Real-Time Policy from 2021:Q4



Note: This figure depicts policy under staggered real-time information about macroeconomic conditions between 2021:Q4 and 2023:Q4. Rate prescriptions are updated every quarter by conditioning the policy projections on the most recent vintage of staff projections at the time (December 2021, March 2022, June 2022, September 2022, December 2022, March 2023, June 2023, September 2023, December 2023) instead of the realized values of macroeconomic variables. The policy is thus set optimally according to the best available forecast at the time but may be suboptimal ex post as unexpected shocks hit the economy. The staggered rule-based policy prescriptions are computed and updated each quarter by finding the policy path that is consistent with the feedback rule of the model while the optimal policy paths are computed by minimizing the central bank intertemporal loss function.

rule of the model (blue diamonds), or optimal policy with hawkish (red diamonds) and benchmark (yellow diamonds) preferences.⁹ The staggered nature of the simulated policy paths reflects changes in the economic outlook, market expectations for the interest rate path, and the implied optimal policy response in each quarter.

The real-time projections suggest that rate hikes could have started in the first half of 2022 and take the policy rate to 4 percent in 2024:Q1. Actual rate hikes instead started in 2022:Q3 and took the policy rate to 4 percent in 2023:Q3. However, unlike simulations with full information about the sequence of shocks hitting the economy, in real time there is no rationale for hiking policy rates aggressively already in 2022:Q1.

⁹Figures A.2 and A.3 (in the appendix) show the quarter-by-quarter projection baselines for interest rates, real GDP growth, and HICP inflation, as well as projections from the corresponding policy counterfactuals.

Overall, the real-time simulations confirm that actual policy fares relatively well against alternative policy conducts. In particular, the actual rate path aligns more closely with real-time optimal policy projections than with the model's estimated interest rate rule. Optimal policy projections in real time from 2021:Q4 indicate a slower tightening compared to what the model's interest rate rule would suggest. However, both real-time rule-based and optimal policy prescriptions still imply an earlier liftoff date, although not as steep as the actual rate path. While policy rates normalization appears to have been delayed in comparison to these real-time projections, the steeper hiking profile has at least partially counterbalanced the initial easing impulse.¹⁰ From 2024 onward, counterfactual policy projections point to policy rates above market expectations, with the exception of the benchmark optimal policy, which closely aligns with the forward curve.

3.5 *Optimized Simple Rules*

Our set of experiments so far has focused on simulating counterfactual rule-based policies and performing counterfactual optimal policy projections. In this next experiment we aim for a middle ground of deriving the optimized simple rule (OSR) (as in Schmitt-Grohe and Uribe 2007), which enables us to find the optimized coefficients for the monetary policy rule in our large-scale NAWM II. Specifically, we search for the set of parameters in Equation (4) that minimize the welfare loss function from Equation (5), that is the variance of (annual) HICP inflation, the output gap, and interest rate volatility, given the other estimated parameters and dynamic properties of the model. In this way, we aim to identify a tractable and transparent policy rule that closely approximates the outcomes of fully optimal policy under commitment, while being more implementable and robust to model uncertainty.

Table 2 provides the coefficients from the optimized simple rules based on the specification of the NAWM rule. The “Benchmark”

¹⁰The ECB initiated the tightening cycle with a 50 basis point hike in July 2022, followed by two 75 basis point hikes in the subsequent policy meetings. Lane (2023) shows that the scale and speed of the adjustments stand out in the history of the monetary union.

Table 2. Optimized Benchmark and Hawkish Policy Rules

Variables		NAWM	Benchmark	Hawkish
ϕ_R	Interest Rate Smoothing	0.9346	0.9993	0.9988
ϕ_π	Response to Inflation	2.9334	183.2040	195.8903
ϕ_Y	Response to Output Gap	0.0320	8.4097	1.4462
$\phi_{\delta\pi}$	Response to Change in Inflation	0.0361	-0.1736	-0.1134
$\phi_{\delta Y}$	Response to Change in Output Gap	0.0920	0.3235	0.1826
Note: The parameters of the optimized simple policy rules are calculated to minimize the variance of annual HICP inflation, the output gap, and the policy rate, given the other estimated parameters and dynamic properties of the model. The “benchmark” optimized rule assigns a weight on output that is half of that on price inflation, whereas the “hawkish” optimized rule assigns a weight on output that is 10 percent of that on inflation.				

column considers the case where the loss function assigns a weight on output that is half of that on price inflation, whereas the “Hawkish” column minimizes the coefficients relative to a loss function which assigns a weight on output that is 10 percent of that on inflation. The “NAWM” column reports the results from the estimated model.

Taking into account the set of optimized coefficients, the hawkish rule assigns a greater weight to inflation stabilization than the estimated NAWM rule, consistent with its definition of the loss function, which places 90 percent of the policy weight on inflation relative to output. The benchmark rule features a sizable coefficient on the output gap, consistent with a more balanced policy stance—specifically, output receives half the weight of inflation in the loss function. In contrast, the hawkish rule assigns a much smaller response to output, reflecting the diminished importance of output stabilization. Both the benchmark and hawkish rules imply high interest rate smoothing parameters (close to 1), reflecting a strong preference for gradualism in policy adjustments. This smoothness helps reduce policy-induced volatility and supports macroeconomic stability by avoiding abrupt changes in the policy stance. Finally, the original policy rule allows for the stabilization over the changes in inflation and the change in the output gap. The small and slightly negative coefficients on the change in inflation in the OSR suggest that the rule modestly dampens overreaction to short-run inflation

movements, acting as a stabilizer against transitory fluctuations. The negative coefficient is interpreted as a corrective force, reducing the amplification of temporary inflation shocks. The optimized coefficients also suggest a moderate response to the change in the output gap, again enhancing the rule's responsiveness to the economic cycle without overreacting to temporary changes. The benchmark rule places a higher weight on this term than the hawkish one, consistent with the former's stronger emphasis on real-side stabilization.


Figure 4 simulates the rule-based policy counterfactual projections in 2022:Q3 based on optimized coefficients and compares their prescriptions against the corresponding (estimated) rule-based projection. As with the rule-based projection, the counterfactual macroeconomic outcomes in the OSR-based projection are computed by resimulating the filtered projection baseline, taking actual policy deviations from the rule as unanticipated monetary policy shocks. The only difference across the set of rules therefore arises from the alternative parameterization of the coefficients.

Our results suggest that the hawkish optimized simple rule performs similarly to the estimated policy rule, prescribing a similar increase in the policy rate and of the same duration for achieving a marginally larger decrease in inflation from baseline in the medium term. Instead, the “benchmark” optimized simple rule—which attaches a higher weight to output gap stabilization—prescribes a more gradual increase in the policy rate, which closes the output gap faster but at the expense of a slower convergence of inflation to target.

3.6 *Stabilization Metrics*

Given the various monetary policy strategies evaluated so far, we provide synthetic indicators of macroeconomic volatility and welfare implied by each strategy. Our focus is on measuring the deviation of relevant variables x from their respective targets $\bar{x}$ over the period 2022:Q1–2026:Q4. For each monetary policy strategy j , our metric of ex post macroeconomic stabilization is defined as the square root of the sum of discounted squared deviations of inflation, the output gap, and quarterly changes to the policy rate from their target levels, 2 percent, 0 percent, and 0 percent, respectively:

Figure 4. NAWM Rule Against Optimized Simple Rules



Note: This figure depicts counterfactual policy projections conditioning on the data and macroeconomic outlook from Eurosystem’s staff projections in December 2023 (solid dark blue lines). Solid light blue lines represent the NAWM rule-based policy projections. Solid red (yellow) lines represent the optimal simple rules projections in 2021:Q4 with optimized coefficients under a minimization of the loss function consistent with hawkish (benchmark) preferences.

$$\mathcal{V}_j = \sqrt{\sum_t^{t+h} \beta^t \mathbb{E}_t (x_{t,j} - \bar{x})^2}. \quad (6)$$

For the evaluation of *real-time* strategies, we take the average over n quarterly macroeconomic and policy projections vintages of the square root of the sum of squared deviations of inflation, output gap, and quarterly changes to the policy rate from their target levels:

$$\tilde{\mathcal{V}}_j = n^{-1} \sum_n \sqrt{\sum_t^{t+h} \beta^t \mathbb{E}_t (x_{t,j} - \bar{x})^2}. \quad (7)$$

Separately, we also report the welfare loss from each strategy, consistent with Equation (5), to assess to what extent various strategies affect consumer utility.

Table 3 illustrates the results. Columns 1–3 present the volatility of inflation, the output gap, and the policy rate, while columns 4 and 5 show the associated welfare loss under alternative calibrations of the loss function.

Starting with the ex post evaluation of the alternative monetary policy strategies, they all improve on the stabilization of inflation relative to actual policy, with the notable exception of the lower-for-longer (LFL) policy where the hiking cycle is initiated in 2024:Q1. Earlier policy tightening and the hawkish optimal policy conduct would deliver the most significant dampening of inflation volatility. Optimal policy with hawkish preferences, in particular, lowers inflation volatility by 32 percent with earlier interest rate liftoff relative to actual policy. However, these inflation stabilization benefits, whether initiated in 2021:Q4 or 2022:Q3, are associated with higher output volatility. Considering the early implementation case, the output volatility is increased from 3- to 10-fold across the optimal and rule-based policy conduct.

In order to assess such alternative management strategies of the output-inflation trade-off, the measures of loss functions (for both benchmark and hawkish preferences) are reported in Table 3, for all policy configurations. Assuming the ex post simulation protocol, it appears that the benchmark optimal policy outperforms actual policy, whatever the loss function is used in the evaluation. However,

Table 3. Macroeconomic Volatility and Welfare with Alternative Strategies

	Volatility			Welfare Loss	
	Inflation	Output	Policy Rate	Benchmark	Hawkish
<i>Ex Post with Full Information</i>					
Actual Policy Rule-Based Policy Early 2022:Q3 2024:Q1 (LFL) Optimal Simple Rules Early, Benchmark Early, Hawkish Optimal Policy Under Commitment Early, Benchmark Early, Hawkish 2022:Q3, Benchmark 2022:Q3, Hawkish	1.00	1.00	1.00	1.00	1.00
	0.79	8.09	1.14	1.15	0.76
	1.00	0.97	1.06	1.01	1.01
	1.24	12.68	0.97	2.71	1.74
	0.90	3.01	0.87	0.87	0.82
	0.83	7.99	0.98	1.17	0.79
	0.87	3.66	0.85	0.85	0.78
	0.68	10.06	1.16	1.25	0.66
Actual Policy Rule-Based Policy Optimal Policy, Benchmark Optimal Policy, Hawkish	0.98	1.90	1.04	0.99	0.97
	0.94	5.07	1.14	1.10	0.95
	<i>Staggered, Real-Time Information</i>				
	1.00	1.00	1.00	1.00	1.00
Actual Policy Rule-Based Policy Optimal Policy, Benchmark Optimal Policy, Hawkish	1.00	1.99	1.16	1.09	1.02
	0.99	1.30	0.96	0.99	0.98
	0.97	1.94	1.02	1.04	0.97
	Note: The figures are provided relative to actual policy over the 2022:Q1–2026:Q4 period: a value above (below) 1 indicates a higher (lower) variance of the policy target variable in the policy counterfactual relative to actual policy. For instance, a value of 0.9 indicates a 10 percent improvement in the stabilization of the policy target variable. Benchmark preferences: $\lambda_\pi = 1$, $\lambda_y = 0.5$, $\lambda_{\Delta i} = 4$. Hawkish preferences: $\lambda_\pi = 1$, $\lambda_y = 0.1$, $\lambda_{\Delta i} = 4$.				

a concomitant implementation of the benchmark optimal policy does not allow significantly outperforming the baseline. The rule-based policy generally delivers worse outcomes, except in the early tightening case for the loss function with the hawkish preferences. Notably, the rule-based policy, when implemented as of 2022:Q3, has a similar performance to the actual policy across all metrics. The hawkish optimal policy only improves the loss function with the hawkish preferences (against which the policy conduct is optimized): this policy would be ranked even worse than the rule-based policy, on the basis the benchmark loss function. Finally, the LFL policy underperforms markedly because looking through the sequence of shocks leads to higher and above-target inflation for at least three more years and to a prolonged recession once the policy tightening eventually materializes in the course of 2024, lasting at least twice as long as the recession in the early-tightening scenario.

Turning to the real-time stabilization metrics between December 2021 and December 2023 (last rows of Table 3), the inflation stabilization or welfare improvements of the optimal policies become smaller. The benchmark optimal policy can still yield some benefits, but they remain below 2 percent for inflation volatility or the loss functions. The hawkish policy only improves welfare if preferences are indeed hawkish.

To summarize, the review of the various stabilization metrics broadly supports the effectiveness of the actual policy conduct since 2021:Q4. However, it is important to acknowledge that the validity of these counterfactual results and their welfare evaluation is inherently tied to the assumptions underlying the model. In particular, the analysis employs a first-order approximation to capture the model's dynamics. Given the sizable shocks during the episodes under investigation, certain nonlinearities may not be fully captured, which could result in less persistent effects or muted responses, thereby understating the welfare losses associated with different strategies. Additionally, the model assumes that longer-term inflation expectations remained anchored throughout the period of investigation. While reasonable in retrospect, this assumption could influence the results by dampening the potential feedback effects on inflation dynamics and hence welfare, especially under conditions of heightened uncertainty (Carrier and Mavromatis 2024).

4. Robustness

As the validity of the findings presented so far is based on key model assumptions, there is considerable uncertainty around several aspects of our analysis. In what follows, we focus on disciplining this uncertainty around three key dimensions: (i) the specification of policymaker preferences in the central bank's loss function, (ii) key parameters of the model that affect the transmission of shocks through the production structure and the financial sector, and (iii) private-sector expectation formation mechanisms.


4.1 *Policymaker Preferences*

The calibration of the loss function utilized so far for optimal policy represents policymaker preferences that range from a balanced dual mandate approach ("benchmark") to a "hawkish" environment. To assess the robustness of our results to policymaker preferences, we extend this range of central bank preferences by computing, for a given relative weight of the price stability coefficient ($\lambda_\pi = 1$), the set of parameters λ_y and $\lambda_{\Delta i}$ which deliver inflation of 2 percent in 2025 or 2026. We do so both for a hiking cycle which is initiated in 2021:Q4 as well as in 2022:Q3.

Figure 5 illustrates the effects from the range of strategies delivering target-consistent inflation in 2025 (yellow area) and 2026 (red dashed line) when initiating the optimal hiking cycle in 2021:Q4 (left column) or 2022:Q3 (right column). These strategies encompass a constellation of policy preferences with respect to macrofinancial volatility considerations while projecting desired inflation outcomes.¹¹ Several target-consistent strategies exist, and therefore

¹¹Figure A.4 in the appendix presents the combinations of preferences (calculated over λ_y and $\lambda_{\Delta i}$) for which an optimal policy is consistent with price stability in the medium term for various policy horizons, 2025 and 2026, respectively. Starting in 2021:Q4 and considering a three-year horizon, there is a positive relationship between the weight assigned to output λ_y and the interest rate smoothing parameter $\lambda_{\Delta i}$, indicating that considerations for inflation-output trade-offs must be associated with more gradual changes to interest rates. Instead, postponing the hiking cycle start results in a shorter policy horizon and requires a more moderate tightening impulse to mitigate the risk of falling below the inflation target in 2025. For price stability over a longer horizon, there is a negative

Figure 5. Optimal Policy Projections with Inflation-Target-Consistent Strategies



Note: This figure depicts the optimal counterfactual policy projections conditioning on the data and macroeconomic outlook from Eurosystem’s staff projections in December 2023 (solid blue lines). The loss function (Equation (5)) is parameterized using the range of parameters contained in the white isocurves of Figure A.4, which deliver target-consistent inflation in 2025 (yellow area) and 2026 (red dashed lines).

relationship between λ_y and $\lambda_{\Delta i}$, implying that a lower weight on output necessitates higher rate smoothing to increase persistence at the peak level and prevent overshooting the inflation target.

price stability in 2025 can be achieved optimally with a lower smoothing parameter, especially when combined with a lower weight on output stabilization. Conversely, additional weight on output stabilization is associated with smoother policy. The resulting optimal path ranges are in the ballpark with those obtained from the simulations under our two main loss functions.


Target-consistent optimal policy projections, computed for a constellation of policy preferences, depend on the policymaker's choice for the liftoff date and the policy-relevant horizon. When the rate liftoff is initiated already in 2021:Q4, optimal policy frontloads rate hikes to slow down economic activity and obtain a lower peak of inflation in 2022 and ensure a rapid disinflation process. If the policymaker targets inflation in 2025, this may result in some undershooting of the target in 2024, while targeting inflation in 2026 would result in a convergence toward the target from above. Delaying the liftoff to 2022:Q3 reduces the scope to affect inflation in 2022 and 2023. To improve over the baseline, optimal policy requires a steeper rate path but a somewhat lower peak compared to the policy initiated in 2021:Q4. If inflation must be at the target in 2026 (and not in 2025), optimal policy is to tighten more to reduce inflation while allowing some undershooting in 2025 before the target is reached from below in 2026.

4.2 *Parameter Uncertainty*

Next, we undertake robustness checks around key parameters of the model that dictate the transmission of shocks through the production structure and the financial sector. For consistency with previous experiments, we focus on optimal policy counterfactuals. Figure 6 illustrates the optimal policy starting in 2021:Q4 under a steeper slope of the price Phillips curve (dashed green line), higher households' portfolio adjustment costs (dotted light blue lines), higher degree of wage indexation (dashed purple lines), and lower investment adjustment costs (dashed red line) contrasting them against our benchmark (solid yellow lines).¹²

¹²These sensitivity exercises also capture instances of possible nonlinearities reflecting, e.g., a state-dependent Phillips curve or nonlinear wage-price effects

Figure 6. Optimal Policy Projections with Parametric Sensitivity



Note: This figure depicts optimal counterfactual policy projections in 2021:Q4 conditioning on the data and macroeconomic outlook from Eurosystem’s staff projections in December 2023 (solid blue lines), in models with alternative parameterization of the slope of the price Phillips curve, the households’ portfolio adjustment costs, the degree of wage indexation, or the strength of investment adjustment costs.

Slope of the Price Phillips Curve. The slope of Phillips curve is one of the parameters of the model that is originally estimated using Bayesian methods and therefore subject to some degree of parametric uncertainty. All else equal, a steeper Phillips curve implies that price inflation is more responsive to changes in economic activity, and therefore a tightening impulse can induce greater disinflation for a given impact on output and employment, and hence a more favorable sacrifice ratio (Tetlow 2022; Ciccarelli, Darracq Pariès, and Priftis 2024).

To quantitatively assess this channel, we set the Calvo parameter ξ_H to 0.7, below the posterior mode of 0.82. In this version of the model, the probability that firms receive permission to optimally reset the price of the output sold in the domestic market is higher. As policy faces a lower sacrifice ratio, a higher policy rate is the optimal policy: under this configuration, higher policy rates would significantly dampen inflationary pressures with a limited additional fall in output, as also shown by Karadi et al. (2024).

Strength of Financial Frictions. The households' portfolio adjustment cost γ_L^h is a key parameter governing the behavior of households and the overall degree of financial frictions in the model. Households hold a portfolio of short-term and long-term government bonds, and any changes to its composition are subject to adjustment costs. The demand for the long-term bond therefore depends on the present value of excess bond returns relative to short-term bonds (i.e., an interest rate premium), which is increasing with an elasticity equal to the inverse of the adjustment cost parameter. As γ_L^h increases, the demand elasticity falls and households respond less flexibly to changes in excess bond returns. Overall, the degree of portfolio adjustment costs amplifies the effects of shocks in the model when the financial intermediary balance sheet constraint tightens and the premium of long-term government bond increases.

We increase the households' portfolio adjustment cost to an upper bound of $\gamma_L^h = 0.30$ instead of the estimated value of $\gamma_L^h = 0.007$. A higher value would hence amplify the response of economic activity to a monetary policy shock, as an increase in the interest rate would lead to a reduction in demand for long-term financing from

and hence a different transmission of shocks in high-inflation regimes (see, e.g., De Santis and Tornese 2023).

the households—by raising excess bond returns by more—which lowers investment, consumption, and output overall. As a result, in an economy with a higher degree of financial frictions, the central bank faces a larger sacrifice ratio and optimal policy consists of a slightly looser stance.

Degree of Wage Indexation. Given the resilience of the labor market during the tightening cycle under investigation and the recent behavior of wages in the euro area (Arce et al. 2024), we next consider sensitivity of optimal counterfactuals to the degree of wage indexation (to inflation) of intermediate goods producers. Wage indexation plays a key role in shaping the transmission of monetary policy by influencing the degree of nominal rigidity in the labor market and the persistence of inflation.

We provide robustness to this mechanism by increasing the degree of wage indexation to $\chi_W = 0.7$ from the benchmark estimated value of $\chi_W = 0.36$. Since higher indexation slows the adjustment of real wages and adds inertia to inflation dynamics, the required peak rate to lower inflation and improve on the output gap is slightly higher and more persistent relative to the benchmark calibration. However, quantitatively, the higher degree of wage indexation produces negligible differences on inflation and the output gap across model environments, as while real wages now adjust more slowly to past inflation, current inflation also reflects the slower reduction in future real marginal costs.

Strength of Investment Adjustment Costs. Among the factors that influence the transmission of a monetary policy, an important role is played by real rigidities. The latter include habit persistence in households' consumption and adjustment costs paid by firms when changing their investment decisions. The purpose of introducing some degree of real rigidities in the model is to reproduce the sluggish response of real variables that is observed in the data following the impact of exogenous shocks. In order to explore the role of real rigidities in the transmission of optimal monetary policy, we lower the strength of investment adjustment costs that intermediate goods firms face, from $\gamma_I = 9$ to $\gamma_I = 1$, such that real rigidities are lowered to the extent possible.

Reducing the real rigidities parameters is akin to removing the sluggishness in the response of output. As a result, the optimal policy counterfactual results in a lower peak rate to stabilize output and

inflation. Notably, despite the lower peak rate, the response of inflation is not significantly different from the baseline parameterization. As output components adjust more rapidly to optimal policy, current inflation also reflects the reduction in future aggregate demand via the Phillips curve.

4.3 Private-Sector Expectation Formation

Optimal policy is implemented under commitment such that policy announcements about future policy affect the contemporaneous behavior of agents via anticipated shocks. In models with rational expectations, these anticipated shocks may lead to implausibly strong effects on macroeconomic variables, the so-called forward-guidance puzzle. To attenuate the propagation of interest rate shocks in the NAWM II, we have so far used the method in Gabaix (2020) by assuming that only 80 percent of private-sector agents in the model are attentive to policy announcements, in line with empirical estimates from Coenen et al. (2022). Specifically, in each period, the impulse responses of the endogenous macroeconomic variables to the anticipated policy shocks are attenuated by the share of private-sector agents that are assumed to be inattentive.¹³

To further check the robustness of our optimal policy results, we next compute optimal policy under alternative expectation formation mechanisms, such as imperfect credibility (Haberis, Harrison, and Waldron 2019), a finite planning horizon (Woodford 2019), and learning as in Cole (2021), following closely the approach in de Groot et al. (2021). Specifically, under imperfect credibility, we let the fraction of agents incorporating the policy shock decaying in the horizon h of the announcement according to α^h with $\alpha = 0.8$. Instead, under a finite planning horizon, agents are assumed to completely dismiss policy shocks anticipated to hit beyond a four-quarter

¹³Given that the degree of dampening of the forward-guidance puzzle has implications for optimal monetary policy (Nakata et al. 2019), we perform additional simulations (not shown here) with different degrees of rational inattention. Fully removing the share of inattentive agents would produce larger effects but appear implausible due to the forward-guidance puzzle. However, our results are broadly robust to varying the degree of attention to policy announcements between the empirically plausible range of 90 percent and 50 percent.

period in the formation of their expectations. Finally, under learning, we assumed that a fraction of agents initially dismiss optimal policy announcements, but this fraction declines over time, rendering agents to progressively understand the central bank's commitment with more certainty. In particular, in the learning case, the degree to which anticipated policy shocks are integrated into expectations is defined by a two-parameter logistic function, where $\beta_1 = 1$ defines the speed of learning about the central bank's commitment to its announcement, and $\beta_2 = 5$ defines the initial beliefs.


The results are summarized in Figure 7,¹⁴ and overall suggest that the way in which bounded rationality of the private sector is introduced affects the aggressiveness of optimal central bank policy both in terms of speed of tightening and the peak rate achieved as well as the persistence of the restrictive stance.¹⁵

On the one hand, under imperfect credibility, unlike under inattention, the share of agents incorporating the policy shocks into their expectations decays exponentially with the policy horizon. As a result, policy announcements further away in the future are therefore even less effective under imperfect credibility than under inattention, which necessitates a lower peak rate, as agents would incrementally observe the policy easing impulse, relative to the baseline. Similarly, under learning, much of the tightening is frontloaded, as policy shocks are largely unanticipated in the short run, while in subsequent periods the degree of anticipation steadily increases so the optimal policy stabilizes at a level somewhat below. These results lend support to Giannoni, Preston, and Eusepi (2020), which argues that the central bank should be less aggressive following inflationary pressures when the private sector deviates from rationality. However, under finite planning horizons, agents completely ignore policy shocks beyond four quarters ahead, implying a more gradual tightening but ultimately higher peak rates. The latter alternative of deviating from rational expectations lends instead support to Gaspar, Smets, and Vestin (2010), Gàti (2023), and Hommes

¹⁴Figure A.5 in the appendix presents the associated optimal policy projections for each type of alternative expectation formation mechanism.

¹⁵Notably, Carrier and Mavromatis (2024) show that when both the interest rate and central bank asset holdings are set optimally, expectations are found to significantly influence the optimal interest rate trajectory.

Figure 7. Optimal Policy Projections with Alternative Expectation Formation Mechanisms



Note: This figure depicts optimal counterfactual policy projections conditioning on the data and macroeconomic outlook from Eurosystem’s staff projections in December 2023 (solid blue lines). Optimal policy paths are computed from 2021:Q4 onward, under commitment, assuming 20 percent inattention of the private sector to future policy shocks (solid yellow lines). The yellow area includes the range of other private-sector expectation formation mechanisms: imperfect credibility, finite planning horizons, and learning.

et al. (2023), who suggest that an optimal monetary policy is tighter following inflationary pressures in a boundedly rational environment.

5. Conclusion

We have assessed the impact of monetary policy on macroeconomic outcomes by focusing on changes to interest rate policy in the post-pandemic environment. We employed the NAWM II, the workhorse DSGE model for policy analysis at the ECB, and confronted it with alternative, counterfactual monetary policy strategies since mid-2021. Our intention was to contrast the impact from actual ECB policy against alternative strategies that differ in their timing, optimality, and reflection of monetary policymaker preferences.

Our findings indicate that the ECB managed to stabilize inflation and output relatively effectively, albeit with some trade-offs. According to the model-based analysis, the actual policy path—starting rate hikes in 2022:Q3—led to a short-lived economic slowdown, with inflation projected to return to target by end-2025 and interest rates peaking at 4 percent. An earlier tightening in 2021:Q4 would have lowered peak inflation to 8 percent, instead of 10 percent, but triggered a more prolonged recession. Conversely, postponing the hike to 2024:Q1 would have supported short-term growth but allowed inflation to remain elevated—above 4 percent until 2026. Simulations of optimal policy suggest that a more aggressive stance could have further reduced inflation volatility, though at the cost of a deeper downturn. Real-time counterfactuals show that actual policy closely tracked the optimal path, effectively balancing inflation control with output stabilization. Most alternative strategies, aside from a delayed response, enhanced price stability but exacerbated output volatility—underscoring the ECB's success in navigating the inflation-growth trade-off.

Finally, it is worth acknowledging that the validity of these counterfactual results depends on key model assumptions. The use of a first-order approximation may not fully capture nonlinearities from sizable shocks, while the assumption of anchored longer-term inflation expectations may reduce feedback effects on inflation. We leave these extensions to further research.

Appendix

A.1 AIT vs. Optimal Policy

Given the hump-shaped nature of our optimal interest rate paths, we explore whether an average inflation targeting (AIT) framework can produce similar effects, as was the implemented policy in the United States during the post-pandemic tightening cycle.

Specifically, we run rule-based policy projections where the central bank sets interest rates according to an AIT rule, by targeting the moving average of the inflation gap. We assess how the standard inflation targeting rule compares to the one-year AIT rule, which is commensurate with targeting year-on-year inflation, as well as two-year AIT, three-year AIT, and four-year AIT. The experiment differs from the optimal policy projections to the extent that the AIT rule is backward looking, as the central bank sets the policy rate according to a systematic response to observed shortfalls of inflation from the target over several quarters, while optimal policy is forward looking and instead consists of minimizing the projected welfare loss.

The AIT specification over q quarters requires the interest rate rule to now take the following form:


$$i_t = \rho i_{t-1} + (1 - \rho) \left(\bar{r} + \bar{\pi} + \phi_{\hat{\pi}} \frac{1}{q} \sum_{j=0}^{q-1} \hat{\pi}_{t-j} + \phi_{\Delta\pi} \Delta\pi_t + \phi_{\hat{y}} \hat{y}_t + \phi_{\Delta\hat{y}} \Delta\hat{y}_t \right) + \eta_t, \quad (\text{A.1})$$

where the inflation gap is now replaced by the average inflation gap over the last q quarters.

Conditioning on the December 2023 baseline, our results suggest that—in hindsight—an AIT rule would prescribe an earlier liftoff from the effective lower bound already in 2021:Q4. The results are similar to the evidence in Crump et al. (2023), which provides a model-based prediction using the FRB/US model by allowing the U.S. Federal Reserve to follow an AIT rule, but are also in line with the optimal policy and (estimated) rule-based prescriptions from our analysis. However, the speed of hiking is more gradual under the AIT

rule relative to optimal policy because AIT takes into account accumulated inflation shortfalls before the inflation surge, while optimal policy is conditional on inflation forecasts. Given that the average inflation rate was below 2 percent before 2021, when inflation started to rise above the target in the second half of 2021, a central bank stabilizing a moving average of the inflation rate tolerates above-target inflation until the average inflation gap closes. As a result, under an AIT rule, the central bank hikes the policy rate once the inflation gap turns positive and maintains it persistently higher until inflation undershoots the target to allow the average inflation gap to close in the medium term. In particular, the gradualism of rate hiking, i.e., the number of quarters required to reach the peak rate, increases with the duration of the AIT window.


Figure A.1. AIT Rule-Based Counterfactuals in 2021:Q4



Note: This figure depicts counterfactual AIT rule-based policy in 2021:Q4 conditioning on the data and macroeconomic outlook from Eurosystem’s staff projections in December 2023 (solid blue lines). Under AIT rules, the central bank targets average inflation over a period between one year and four years.


A.2 Real-Time Counterfactuals

Figure A.2. Real-Time Forecasts between December 2021 and December 2023



Note: This figure depicts counterfactual rule-based policy (right) conditioning on quarterly real-time ECB/Eurosystem’s staff projections (left) between 2021:Q4 and 2023:Q4. The policy projections are computed and updated each quarter by finding the policy path consistent with the feedback rule of the model.


Figure A.3. Real-Time Forecasts between December 2021 and December 2023



Note: This figure depicts counterfactual optimal policy projections under alternative parameterizations of the central bank function conditioning on quarterly real-time ECB/Eurosystem’s staff projections (left) between 2021:Q4 and 2023:Q4.

A.3 Policy Preferences

Figure A.4 depicts projected annual HICP inflation in 2025 (top row, panels I and II) and 2026 (bottom row, panels III and IV) under counterfactual optimal policy simulations that start in 2021:Q4 (left) or 2022:Q3 (right) conditional on the December 2023 macroeconomic projection baseline. Each point on each grid represents a certain parameter combination of λ_y and $\lambda_{\Delta i}$.




Note: This figure depicts projected annual HICP inflation in 2025 (top row) and 2026 (bottom row) under optimal policy counterfactual simulations that start in 2021:Q4 (left) or 2022:Q3 (right) conditioning on the data and macroeconomic outlook from Eurosystem’s staff projections in December 2023. Each panel is constructed using alternative combinations of parameter values for λ_y and $\lambda_{\Delta i}$, the weights of output and interest rate smoothing relative to price stability in the loss function. The pairs of parameters delivering 2 percent inflation in 2025 and 2026 are interpolated linearly.

If policymakers aim to achieve price stability in 2025 (panels I and II), λ_y and $\lambda_{\Delta i}$ are positively associated. For the combination of shocks affecting the economy over that period, a more gradual steering of the interest rate due to output considerations must be offset by greater persistence of rate path. When optimal policy is initiated in 2022:Q3, the policy horizon shrinks by three quarters. Lags inherent to monetary policy transmission entail a more moderate tightening impulse overall to stabilize inflation at the target in 2025: all else equal, a higher weight on output is compatible with price stability, i.e., the isocurve shifts to the right. A later start therefore implies reduced macrofinancial volatility compared to an earlier start. If policymakers instead aim to achieve price stability in 2026 (panels III and IV), λ_y and $\lambda_{\Delta i}$ are weakly associated. Target-consistent preferences for a policy initiated in 2021:Q4 require a weight on output between 0.2 and 0.3 for any given weight on interest rate smoothing. Instead, λ_y and $\lambda_{\Delta i}$ would be negatively associated when policy is initiated later, implying that higher output considerations must be compensated for with lower smoothing and vice versa.

A.4 Sensitivity to Expectations

Figure A.5. Optimal Policy Projections with Alternative Expectation Formation Mechanisms



Note: This figure depicts counterfactual optimal policy projections conditioning on the data and macroeconomic outlook from Eurosystem’s staff projections in December 2023 (solid blue lines). Optimal policy paths are computed from 2021:Q4 onward, under commitment, assuming inattention of the private sector to future policy shocks (yellow), imperfect credibility (green), four-quarter finite planning horizons (blue), and learning (red).

References

- Adjemian, S., H. Bastani, M. Juillard, F. Karamé, F. Mihoubi, W. Mutschler, J. Pfeifer, M. Ratto, N. Rion, and S. Villemot. 2022. “Dynare: Reference Manual Version 5.” Dynare Working Paper No. 72, CEPREMAP.
- Adjemian, S., M. Darracq Pariès, and S. Moyen. 2007. “Optimal Monetary Policy in an Estimated DSGE for the Euro Area.” Working Paper No. 803, European Central Bank.
- Adolfson, M., S. Laséen, J. Lindé, and L. E. O. Svensson. 2011. “Optimal Monetary Policy in an Operational Medium-Sized DSGE Model.” *Journal of Money, Credit and Banking* 43 (4): 1287–1331.
- Altavilla, C., and M. Ciccarelli. 2011. “Monetary Policy Analysis in Real-Time. Vintage Combination from a Real-Time Dataset.” Working Paper No. 274, Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy.
- Arce, M. Ciccarelli, C. Montes-Galdón, and A. Kornprobst. 2024. “What Caused the Euro Area Post-pandemic Inflation?” Occasional Paper No. 343, European Central Bank.
- Barnichon, R., and G. Mesters. 2023. “A Sufficient Statistics Approach for Macro Policy.” *American Economic Review* 113 (11): 2809–45.
- Brand, C., M. Obstbaum, G. Coenen, D. Sondermann, R. Lydon, V. Ajevskis, F. Hammermann, et al. 2021. “Employment and the Conduct of Monetary Policy in the Euro Area.” Occasional Paper No. 275, European Central Bank.
- Calvo, G. A. 1983. “Staggered Prices in a Utility-Maximizing Framework.” *Journal of Monetary Economics* 12 (3): 383–98.
- Carlstrom, C. T., T. S. Fuerst, and M. Paustian. 2017. “Targeting Long Rates in a Model with Segmented Markets.” *American Economic Journal: Macroeconomics* 9 (1): 205–42.
- Carrier, A., and K. Mavromatis. 2024. “Optimal Normalization Policy under Behavioral Expectations.” Working Paper No. 800/2024, De Nederlandsche Bank.
- Cecion, M., G. Coenen, R. Gerke, H. Le Bihan, R. Motto, P. Aguilar, V. Ajevskis, et al. 2021. “The ECB’s Price Stability Framework: Past Experience, and Current and Future Challenges.” Occasional Paper No. 269, European Central Bank.

- Ciccarelli, M., M. Darracq Pariès, and R. Priftis. 2024. “ECB Macroeconometric Models for Forecasting and Policy Analysis: Development, Current Practices and Prospective Challenges.” Occasional Paper No. 344, European Central Bank.
- Coenen, G., P. Karadi, S. Schmidt, and A. Warne. 2018. “The New Area-Wide Model II: An Extended Version of the ECB’s Micro-founded Model for Forecasting and Policy Analysis with a Financial Sector.” Working Paper No. 2200, European Central Bank.
- Coenen, G., C. Montes-Galdón, A. Saint Guilhem, J. Hutchinson, and R. Motto. 2022. “Rate Forward Guidance in an Environment of Large Central Bank Balance Sheets: A Eurosystem Stock-taking Assessment.” Occasional Paper No. 290, European Central Bank.
- Cole, S. J. 2021. “Learning and the Effectiveness of Central Bank Forward Guidance.” *Journal of Money, Credit and Banking* 53 (1): 157–200.
- Crump, R., M. Del Negro, K. Dogra, P. Gundam, D. Lee, R. Nallamotu, and B. Pacula. 2023. “The New York Fed DSGE Model Perspective on the Lagged Effect of Monetary Policy.” Federal Reserve Bank of New York, *Liberty Street Economics* (blog), November 21.
- Darracq Pariès, M., and M. Kühl. 2016. “The Optimal Conduct of Central Bank Asset Purchases.” Working Paper No. 1973, European Central Bank.
- De Fiore, F., B. Mojon, D. Rees, and D. Sandri. 2023. “Monetary Policy Frameworks Away from the ELB.” BIS Working Paper No. 1156.
- de Groot, O., F. Mazelis, R. Motto, and A. Ristiniemi. 2021. “A Toolkit for Computing Constrained Optimal Policy Projections (COPPs).” Working Paper No. 2555, European Central Bank.
- De Santis, R. A., and T. Tornese. 2023. “Energy Supply Shocks’ Nonlinearities on Output and Prices.” Working Paper No. 2834, European Central Bank.
- Debortoli, D., J. Kim, J. Lindé, and R. Nunes. 2019. “Designing a Simple Loss Function for Central Banks: Does a Dual Mandate Make Sense?” *Economic Journal* 129 (621): 2010–38.

- Dengler, T., R. Gerke, S. Giesen, D. Kienzler, J. Röttger, A. Scheer, and J. Wacks. 2024. “A Primer on Optimal Policy Projections.” Technical Report.
- European Central Bank. 2016. “A Guide to the Eurosystem/ECB Staff Macroeconomic Projection Exercises.” (July).
- Gabaix, X. 2020. “A Behavioral New Keynesian Model.” *American Economic Review* 110 (8): 2271–2327.
- Gaspar, V., F. Smets, and D. Vestin. 2010. “Inflation Expectations, Adaptive Learning and Optimal Monetary Policy.” In *Handbook of Monetary Economics*, Vol. 3, ed. B. M. Friedman and M. Woodford, 1055–95 (chapter 19). Elsevier.
- Gàti, L. 2023. “Monetary Policy & Anchored Expectations—An Endogenous Gain Learning Model.” *Journal of Monetary Economics* 140 (Supplement): S37–S47.
- Gertler, M., and P. Karadi. 2011. “A Model of Unconventional Monetary Policy.” *Journal of Monetary Economics* 58 (1): 17–34.
- . 2013. “QE 1 vs. 2 vs. 3. . . : A Framework for Analyzing Large-Scale Asset Purchases as a Monetary Policy Tool.” *International Journal of Central Banking* 9 (Supplement 1): 5–53.
- Giannoni, M., B. Preston, and S. Eusepi. 2020. “On the Limits of Monetary Policy.” Technical Report.
- Giannoni, M. P., and M. Woodford. 2003. “Optimal Interest-Rate Rules: I. General Theory.” NBER Working Paper No. 9419.
- Haberis, A., R. Harrison, and M. Waldron. 2019. “Uncertain Policy Promises.” *European Economic Review* 111 (January): 459–74.
- Harrison, R., and M. Waldron. 2021. “Optimal Policy with Occasionally Binding Constraints: Piecewise Linear Solution Methods.” Working Paper No. 911, Bank of England.
- Hommes, C., K. Mavromatis, T. Özden, and M. Zhu. 2023. “Behavioral Learning Equilibria in New Keynesian Models.” *Quantitative Economics* 14 (4): 1401–45.
- Karadi, P., A. Nakov, G. N. Barrau, E. Pasten, and D. Thaler. 2024. “Strike While the Iron is Hot: Optimal Monetary Policy with a Nonlinear Phillips Curve.” BIS Working Paper No. 1203.
- Lane, P. R. 2023. “The Euro Area Hiking Cycle: An Interim Assessment.” Dow Lecture, National Institute of Economic and Social Research, London, February 16.

- Mazelis, F., R. Motto, and A. Ristiniemi. 2023. "Monetary Policy Strategies for the Euro Area: Optimal Rules in the Presence of the ELB." Working Paper No. 2797, European Central Bank.
- McKay, A., and C. K. Wolf. 2023. "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" *Econometrica* 91 (5): 1695–1725.
- Molodtsova, T., A. Nikolsko-Rzhevskyy, and D. H. Papell. 2008. "Taylor Rules with Real-time Data: A Tale of Two Countries and One Exchange Rate." *Journal of Monetary Economics* 55 (Supplement): S63–S79.
- Nakata, T., R. Ogaki, S. Schmidt, and P. Yoo. 2019. "Attenuating the Forward Guidance Puzzle: Implications for Optimal Monetary Policy." *Journal of Economic Dynamics and Control* 105 (August): 90–106.
- Orphanides, A. 2001. "Monetary Policy Rules Based on Real-Time Data." *American Economic Review* 91 (4): 964–85.
- Schmitt-Grohe, S., and M. Uribe. 2007. "Optimal Simple and Implementable Monetary and Fiscal Rules." *Journal of Monetary Economics* 54 (6): 1702–25.
- Svensson, L. E. O. 2005. "Monetary Policy with Judgment: Forecast Targeting." *International Journal of Central Banking* 1 (1): 1–54.
- Svensson, L. E. O., and R. J. Tetlow. 2005. "Optimal Policy Projections." *International Journal of Central Banking* 1 (3): 177–207.
- Svensson, L. E. O., and M. Woodford. 2003. "Implementing Optimal Policy through Inflation-Forecast Targeting." NBER Working Paper No. 9747.
- 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.
- Woodford, M. 2001. "Fiscal Requirements for Price Stability." *Journal of Money, Credit and Banking* 33 (3): 669–728.
- . 2010. "Robustly Optimal Monetary Policy with Near-Rational Expectations." *American Economic Review* 100 (1): 274–303.
- . 2019. "Monetary Policy Analysis When Planning Horizons Are Finite." In *NBER Macroeconomics Annual 2018*, Vol. 33, ed. M. Eichenbaum and J. A. Parker, 1–50. University of Chicago Press.

ECB Euro Liquidity Lines*

Silvia Albrizio,^a Iván Kataryniuk,^b Luis Molina,^b
and Jan Schäfer^c

^aInternational Monetary Fund

^bBanco de España

^cCEMFI

Central bank liquidity lines have gained importance as a cross-currency liquidity management tool with the intention to prevent threats to financial stability. We provide a complete timeline of ECB liquidity line announcements and quantify their signaling effect. The premium to borrow euros in FX markets is estimated to decrease by 51 basis points upon announcement. Moreover, we test for spillbacks and find that bank stock prices increase by around 2.1 percent in euro-area countries highly exposed via banking linkages to currencies targeted by liquidity lines. The mechanism behind this novel result is illustrated in a stylized bank default model.

JEL Codes: E44, E58, F33, G15.

*We are grateful to two anonymous referees, Ricardo Reis, Iñaki Aldasoro, Ambrogio Cesa Bianchi, Dmitry Khametshin, Javier J. Pérez, Antonio Millaruelo, Maurizio Habib, David Lodge, Arnaud Mehl, Massimo Ferrari, Stephanie Titzck, Roland Beck, Jorge Ivan Canales Kriljenko, Vina Nguyen, Rafael Repullo, the members of the ESCB Workstream on Globalisation, and audiences at the BIS/BoE/ECB/IMF Spillover Conference, the Sveriges Riksbank Conference, the ECB IPA seminar, and the BdE internal seminar for the useful comments. Thanks to Elisa Sanchez, Enrique Esteban, and Mamen Castillo for their availability and data help, and Carlos Moreno Perez, Begoña Lara, and Najat Bazah Lamchanna for the research assistance. The views expressed herein are those of the authors and should not be attributed to the IMF, its Executive Board, or its management, and neither to the Banco de España nor the Eurosystem. Schäfer acknowledges financial support from Grant PRE2022-101467 funded by MCIN/AEI/ 10.13039/501100011033 and by “ESF +”, the Maria de Maeztu Unit of Excellence CEMFI MDM-2016-0684 and CEMFI CEX2020-001104-M, funded by MCIN/AEI/10.13039/501100011033, and CEMFI. Author emails: SALbrizio@imf.org, ivan.kataryniuk@bde.es, lmolina@bde.es, jan.schaefer@cemfi.edu.es.

1. Introduction

Central bank swap and repo lines have been used extensively in the last two decades to provide foreign-currency liquidity across jurisdictions. These liquidity lines have two components, called legs. The first consists of an agreement between a source central bank and a recipient one. When a swap line is active, the recipient central bank can access the source central bank's currency in exchange for its domestic currency at the spot exchange rate, up to a maximum amount, and at a fixed interest rate, which is below the market rate. At maturity, the same amount of money is exchanged among the two counterparties at the same fixed spot exchange rate.¹ In this way, the recipient central bank can inject liquidity into the domestic market via repo lines with the domestic financial institutions—which is the second leg of the line. Ultimately, these liquidity injections prevent market pressure on the domestic currency, if currency needs were to be met by private agents, or avoid exhausting domestic reserves, if needs were to be met by reserves' sell-off (Aizenman, Jinjarak, and Park 2011).

In this *do ut des* agreement, on the one hand, the recipient central bank can support the liquidity needs of its domestic banking system. On the other hand, by reducing foreign liquidity shortage, the source central bank prevents negative spillovers in the form of financial instability. Note that the source central bank is exempted from bearing the risk of default of the institutions accessing the credit in the recipient country, which is entirely borne by the recipient central bank, reflecting its comparative advantage in monitoring these institutions. The transaction does not involve exchange rate risk for any of the two counterparts, since the forward exchange rate is fixed. Finally, the agreement can be arranged as a swap or repo line.² In the latter case, the recipient central bank has to pledge assets denominated in the source-country currency as collateral to the source central bank to have access to the currency.

¹Although these arrangements are referred to as central bank swap lines, they have different features with respect to foreign exchange (FX) swaps in private markets. De facto, the former provides a collateralized loan of the source central bank's currency.

²Bahaj and Reis (2023) provide an in-depth discussion of the institutional framework and mechanisms of the central bank liquidity lines.

The objective of this paper is to analyze the effect of European Central Bank (ECB) euro liquidity lines. The euro represents the second most important currency in the international monetary system.³ Just during the first month of the COVID-19 pandemic, the ECB had established eight new euro liquidity facilities and most of them have been extended twice, as of September 2022. In spite of this, the ECB liquidity line network and the effectiveness of these tools are, to the best of our knowledge, mostly unknown. We provide a description of the deployment of this tool by the ECB, creating a timeline of swap and repo line announcements based on ECB press releases. We test for the signaling effect of the ECB euro liquidity lines in reducing liquidity tensions and generating positive spillbacks to euro-area (EA) countries. We use announcement dates instead of the actual activation of the line for two reasons: first, the announcement is publicly available, and therefore any signaling effect will appear on that date. Second, if these lines serve as a prevention tool, the announcement should be sufficient to provide confidence in the functioning of the foreign exchange swap market and increase liquidity. Schnabel and Panetta (2020) underline that these liquidity arrangements do not need to be actually used to be effective.

As shown in Bahaj and Reis (2022), by providing an outside option to the FX market, central bank liquidity lines put a ceiling on deviations from covered interest parity (CIP), lowering recipient-country agents' funding costs. Moreover, by the expectations theory of the term structure, this argument also holds in the case of the mere announcement of a liquidity line, which ultimately caps the price of forward agreements. To test empirically the effectiveness of ECB euro liquidity lines, we follow Bahaj and Reis (2022) and Cetorelli, Goldberg, and Ravazzolo (2020). We measure the euro funding costs in FX markets using the daily deviations of the covered interest parity before and after the announcement of a euro liquidity line on currencies targeted by the line versus a control, non-targeted group. Similar to Bahaj and Reis (2022), our identification relies on a difference-in-differences (DID) strategy. Considering the lines the ECB established with small European countries, we find

³ According to the ECB, the euro represents around 36 percent of the share of global payments at the end of 2022 and around 20 percent of the share in global holdings of foreign reserves.

that the mere announcement of euro swap and repo lines reduces the euro funding cost in foreign exchange swap markets by 51 basis points. This effect corresponds to about 40 percent of one standard deviation of the residuals.

Turning to the spillback effects, we rationalize the mechanism of transmission in a stylized model. The announcement of a central bank liquidity line decreases the default probability of recipient-country financial intermediaries by reducing their euro refinancing cost. This increases the expected payoff of the source-country credit to the recipient bank, increasing the source-country bank's value and eventually its stock market price. We test the prediction of the model leveraging on the heterogeneity across euro-area countries. Banks in different EA jurisdictions can be affected heterogeneously depending on their exposures to foreign markets via the domestic banking sector, proxied by the share of cross-border claims of EA banks toward countries that receive liquidity lines. By introducing the exposure of the banking sector in EA countries in a DID framework, we estimate the differential effect of these announcements on the change in euro-area banks' stock prices. Consistent with our model, we find that EA countries with the most exposed banking sectors benefit the most from the announcement of the lines since they experience a relative increase in their stock prices of about 2.1 percent in a four-day window around the announcement. In other words, more exposed banks see their market valuation increase after the announcement. Overall, the analysis suggests that the signaling effect of central bank liquidity facilities is effective in generating a positive direct effect on foreign FX markets as well as spillovers on the source country.

This paper is connected to the economic literature looking at the beneficial effects of swap lines. Several papers have studied how central bank liquidity lines have lubricated both money markets and foreign exchange swap markets in the Global Financial Crisis (GFC) (Carré and Le Maux 2020; Obstfeld, Shambaugh, and Taylor 2009) as well as in the more recent COVID-19 crisis (Aldasoro et al. 2020). Consistent with the central role of the USD in the global financial markets, most of these arrangements have provided liquidity in this currency. Indeed, Federal Reserve (Fed) USD swap lines played an effective lender-of-last-resort function in FX markets by putting a ceiling on deviations from the CIP (Baba and Packer 2009; Bahaj

and Reis 2022; Moessner and Allen 2013, among others). This mechanism also worked during the COVID-19 crisis (Bahaj and Reis 2020; Cetorelli, Goldberg, and Ravazzolo 2020). Aizenman, Ito, and Pasricha (2022) show that trade and banking linkages with the U.S. are positively associated with access to Fed swap and repo lines during the COVID-19 crisis. Fed liquidity facilities announcements during 2020 have led to an appreciation of the recipient countries' currency with respect to the USD, while dollar auctions by major central banks have expansionary effects on other economies. Using micro-level data on FX forward and swap transactions, Ferrara et al. (2022) find that dealers that draw on swap lines reduce their demand for dollars at the forward leg in the FX market and increase their supply of dollars to non-financial institutions, improving market liquidity. Moreover, Bahaj and Reis (2022) find that Fed liquidity lines have a positive spillback effect on the source country by encouraging capital inflows into USD-denominated assets. In this way, while swap and repo lines can be thought of as a by-product of globalization in financial markets since they responded to the liquidity needs of an integrated global financial system, they also proved to be useful to reinforce the international role of the source currency in the international monetary system as well as in international trade. People's Bank of China's 38 swap lines in less than a decade are a clear example of alternative use of such tools (Bahaj and Reis 2026).

The paper is structured as follows: Section 2 provides an overview of the ECB repo and swap lines network across time; Section 3 presents the data used in the analysis; Section 4 presents the empirical analysis of the direct effect of ECB liquidity line on the euro funding costs in FX markets; Section 5 introduces our stylized model, explaining one of the mechanisms behind the positive spillback effect of these facilities, and tests it empirically.

2. ECB Liquidity Lines

Historically, central banks' liquidity lines have been used for three main objectives: (i) defend a peg system (Bordo, Humpage, and Schwartz 2015) or, more broadly, fund FX interventions (Bahaj and Reis 2023), (ii) offer a financial stability tool, which in the case of the Fed becomes a global liquidity backstop (Bahaj and Reis 2022,

among others), (iii) enhance the international use of the domestic currency (Bahaj and Reis 2026).

The first objective led to the establishment in 1962 of the Fed reciprocal currency agreements, i.e., swap lines, first with the Bank of France and by the end of the same year with nine other key central banks. Under the Bretton-Woods system, the Fed intervened in forward foreign exchange markets to reestablish confidence in the USD and to defend its gold peg. At the end of the 1990s, this tool was discontinued. The global swap network regained importance as a cooperation tool across central banks only following the September 11 terrorist attack and more extensively in 2007 and 2008 with the Global Financial Crisis (GFC).⁴ In this context, central banks' liquidity agreements transformed the Fed into the global lender of last resort, limiting fire sales and helping contain the risk of market contagion.

The G10 central banks contributed to this coordinated effort to expand the USD liquidity provision capacity.⁵ In 2011, the network agreed on extending the scope of these swap lines to provide liquidity in each jurisdiction in any of their currencies. By 2013 these temporary bilateral liquidity swap arrangements were converted into a standing agreement with an unlimited withdrawal amount.

Not only did the ECB participate in such a network⁶ and establish swap lines with the People's Bank of China, the Bank of England, and the Swiss National Bank to provide foreign-denominated liquidity, but between October 2008 and August 2022 it established 14 new swap and repo lines to provide euro liquidity. The ECB has continued to renew and expand its euro facilities particularly during the COVID crisis. Euro lines have been extended both in terms of the time frame and the volume, and their pricing conditions changed.

⁴The evolution of the USD swap lines has been described in Goldberg, Kennedy, and Miu (2011), McCauley and Schenk (2020), and Moessner and Allen (2010), among others.

⁵Participating central banks were the Federal Reserve, the Bank of Canada, the Bank of England, the Bank of Japan, the European Central Bank, and the Swiss National Bank.

⁶Since 2011 the agreement also enabled central banks of the network of temporary bilateral liquidity swap arrangements to provide liquidity operations, should they be needed, in Japanese yen, British pound sterling, Swiss francs, Canadian dollars, and euros, in addition to U.S. dollars.

Counterparts were mostly EU countries outside the EA (Bulgaria, Croatia, Czech Republic, Denmark, Hungary, Latvia, Poland, Romania, and Sweden); however, since March 20, 2020, liquidity lines agreements were signed also with non-EU countries (Serbia, San Marino, Albania, and Republic of North Macedonia).⁷ In addition, on June 25, 2020, as part of the pandemic-related crisis response, the ECB announced the establishment of the EUREP, a new and initially temporary Eurosystem repo facility to provide euro liquidity to a broader set of non-euro-area central banks.

Figure 1 provides a complete picture of all ECB liquidity lines announcements until August 2022.⁸ The timeline comprises announcements related to new agreements, time and volume extensions, and changes in the conditions. The latter may regard changes within the first leg of the line, i.e., on the agreement between the two central banks. An example is the March 15, 2020 announcement, which decreased the rate on the standing U.S. dollar liquidity swap arrangements by 25 basis points. However, they may also refer to changes in the second leg of the line, which refers to the open market operations that the recipient central bank put in place to provide liquidity to the domestic financial system. For example, it may entail a change in the frequency or the maturity of the liquidity-providing operations.⁹

Although most of the lines within the network of the bilateral agreements among the G10 countries' central banks are reciprocal in nature, often these lines have been explicitly motivated by the need of providing USD liquidity. To capture the actual intention of the ECB in providing liquidity in euros versus other currencies, Figure 1 distinguishes between euro and non-euro lines based on the

⁷Previous to October 2008, in 2007, the ECB established its first euro swap line with the central bank of Sweden. However, this event is not included in the analysis since the goal of this paper is to estimate the signaling effect of these lines and the latter has been announced in the ECB press releases much later than the actual decision.

⁸The timeline does not include the EUREP facility since the respective country-specific announcements are not public. This remains the case under the ECB's permanent framework for euro liquidity lines (January 2024) and the standing EUREP facility (February 2026), under which the ECB discloses only aggregate weekly drawings across repo lines.


⁹The timeline does not include discontinuations of central bank lines or maturities of liquidity-providing operations with the domestic financial system.

information reported in the ECB press releases. Announcements related to the former are reported below the timeline, while announcements related to the latter are above the timeline. As the chart shows, euro liquidity lines are mainly clustered around crisis episodes, suggesting the ECB liquidity lines have been mainly used as a liquidity backstop. However, some of these have also been used in the context of the exchange rate mechanism (ERM II), a peg system to the euro, which is a prerequisite for any EU member state to join the euro area—as in the case of Bulgaria on April 22, 2020. Differently from the People’s Bank of China, the ECB has not leveraged on euro liquidity facilities to provide incentives for the internationalization of its currency. Bahaj and Reis (2026) show that the creation of People’s Bank of China RMB swap lines between 2009 and 2018 triggered the jump-start of RMB’s international status. In fact, it increased the probability of using the RMB and its share in the country’s payments. Although we don’t formally test whether the announcements of the euro liquidity lines have increased the probability of using the euro for sales invoicing, there are two main factors that bring us to discard that the ECB has intended to use the liquidity lines for the internationalization of the euro, so far. First, the network is very much geographically concentrated, either in EU countries or countries in close proximity that would use the euro for some of their transactions anyway.¹⁰ Second, the line announcements are concentrated around episodes of actual or potential market turbulence or crisis. Appendix A.3 reports the full timeline of ECB liquidity facilities, specifying additional characteristics of each line announcement.

Figure 2, panel A, shows the total number of announcements by year conditional on the type of facility, swap or repo. Repo facilities require adequate collateral in euro-denominated assets and feature a higher lending rate compared to swap lines. Figure 2, panel B, reports the yearly number of announcements of lines explicitly motivated by the need of providing euro liquidity versus other currencies, such as the U.S. dollar or Swiss franc. Taken together, these two panels underline a shift in the role of the ECB. Initially, the lines served mainly as facilities to enhance USD liquidity in

¹⁰As mentioned, we do not consider the EUREP facility.

Figure 2. ECB Liquidity Facility Announcements



Note: The left panel shows the breakdown of the total number of announcements with respect to the type of facilities. Blue bars denote swap lines, while orange ones refer to repos. The right panel distinguishes between euro, in blue, and non-euro provisions, in orange, according to the explicit motivation provided in the ECB press releases.

the EA via swap lines. During the Global Financial Crisis, both ECB euro and non-euro lines increased predominantly due to the increasing number of swap line extensions with G10 countries. However, the pandemic triggered increasing ECB interventions as regional lender of last resort providing euro-denominated loans to European countries outside the euro area, mainly via repos.

With respect to the conditions, except for the swap lines within the G10 network, most of the facilities were temporary during our sample period, with a maximum amount of allotment at an interest rate specified in each agreement.

Finally, with respect to the activation of the lines, the available information is scarce. Based on recipient countries' sources of information, at least the liquidity lines with the Magyar Nemzeti Bank, the Sveriges Riksbank, and the Bank of England have already been used, although in small amounts.

3. Data

For the purpose of this analysis, we consider all euro liquidity lines reported in the ECB press releases.¹¹ We exclude (i) countries that are targeted by the lines but use the euro as the main currency, such as San Marino; (ii) the G10 network, due to the *de iure* reciprocal nature of the swap lines, to avoid confounding the effect of a euro liquidity line with the impact of USD or other currencies provisions; (iii) the agreements with the National Bank of the Republic of North Macedonia, the National Bank of Romania, and the Bank of Albania, due to data limitation for the construction of the dependent variable; (iv) agreements that were not publicly announced or were not included in the ECB press releases, as is the case for the December 2007 agreement with Sveriges Riksbank and the November 2008 agreement with the Bank of Latvia, since the signaling effect might be diluted. The final sample then comprises 12 liquidity line announcements with seven counterparts whose currencies have been targeted by the lines: Bulgaria, Denmark, Croatia, Hungary, Poland, Serbia, and Sweden. The liquidity lines in our sample were established between October 2008 and March 2022. None of the liquidity lines established prior to October 2008 fit our selection

¹¹The timeline is available at ECB liquidity lines - Timeline.

criteria, due to lack of public announcement (December 2007, Sweden), or reciprocity within the G10 network (all other events prior to October 2008).

The included events are the following:

- October 16, 2008: Repo agreement with the Magyar Nemzeti Bank (HU) to support the bank's instruments of euro liquidity provision.
- October 27, 2008: Swap agreement with the Danmarks Nationalbank (DK) to improve liquidity in euro short-term markets.
- November 21, 2008: Repo agreement with the National Bank of Poland (PL) to support the NBP's instruments of euro liquidity provision.
- June 10, 2009: Activation of the swap agreement signed with the Sveriges Riksbank (SE) on December 20, 2007 with the aim of facilitating the functioning of financial markets and providing euro liquidity to the latter if needed.
- March 20, 2020: Reactivation of the swap line with Danmarks Nationalbank (DK) to provide euro liquidity to Danish financial institutions.
- April 15, 2020: Precautionary swap agreement with the Central Bank of Croatia (HR) to provide euro liquidity to Croatian financial institutions in order to address possible market dysfunction.
- April 22, 2020: Precautionary swap agreement with the Bulgarian National Bank (BG) to provide euro liquidity.
- July 17, 2020: Repo agreement with the National Bank of Serbia (RS) to provide euro liquidity to Serbian financial institutions to address possible euro liquidity needs in the presence of market dysfunctions due to the COVID-19 shock.
- July 23, 2020: Repo agreement with the Magyar Nemzeti Bank (HU) to provide euro liquidity to Hungarian financial institutions to address possible euro liquidity needs in the presence of market dysfunctions due to the COVID-19 shock.
- August 28, 2020: Extension of the repo facility with the Central Bank of Croatia (HR) to provide euro liquidity to financial institutions in the two countries via their respective national central banks to address possible euro liquidity needs in the presence of market dysfunctions due to the COVID-19 shock.

- February 4, 2021: Extension of the temporary swap and repo lines with non-euro-area central banks (HR, RS, HU) to address possible euro liquidity needs in non-euro-area countries in presence of market dysfunctions due to COVID-16.
- March 28, 2022: Swap agreement with Narodowy Bank Polski (PL), and extension of temporary repo lines with central banks of Hungary (HU) to prevent spillover effects in euro-area financial markets and economies that might adversely affect the smooth transmission of the ECB's monetary policy in the context of heightened geopolitical tensions triggered by the Russian invasion of Ukraine.

Finally, as a robustness check, we add two countries whose currencies have never been targeted by ECB liquidity lines: Norway and Iceland. They represent an interesting robustness check because, although both are small countries, close EU partners, and have experienced financial distress, they have not been targeted by the lines.

To test the effectiveness of the ECB euro liquidity line, we consider the change in the euro funding cost in FX markets. In a frictionless FX market, the covered interest parity holds and the implied euro interest rate in the FX market equals the euro money market interest rate. If the CIP does not hold, the FX swap basis spread provides a measure of the premium paid by foreign agents to borrow euros for a specified time period in the FX market compared to the euro money market. In other words, a positive basis represents relatively high costs for euro funding in the FX market. Following Bahaj and Reis (2022), the euro basis is defined as

$$B_t = \ln(F_t) - \ln(S_t) - (r_t - r_t^*), \quad (1)$$

where F_t is the one-week market forward rate of the euro against the rest of the currencies, S_t is the equivalent spot rate, r_t stands for the interest rate of the euro deposits, and r_t^* is the interest rate for deposits in each of the foreign currencies considered.¹² We

¹²Under covered interest parity, the no-arbitrage condition $\frac{F_t}{S_t} = \frac{1+r_t}{1+r_t^*}$ holds. Equation (1) follows from taking logs and using the approximation $\ln(1+r) \approx r$, valid for small r .

use one-week euro currency deposits when possible, and build back some series for some countries using the equivalent interbank interest rate. All data are obtained from Refinitiv databases on a daily frequency. We prefer to use OIS rates, but due to data availability constraints in some cases, we use EURIBOR (euro-area interbank offered rate) rates instead. Table A.6 in Appendix A.4 shows the main descriptive statistics for CIP deviations. Furthermore, since recipient-country idiosyncratic factors may drive the FX basis in times of financial turbulence, we purge the basis from country-specific factors closely related to the occurrence of financial crises, such as sovereign defaults, and banking runs, or currency crashes.

In other words, in the econometric specification, our main dependent variable is defined as the residuals res_{jt} of country-by-country regressions of the basis on country-specific characteristics collected in the vector $\Omega_{j,t}$:

$$basis_{jt} = \alpha_j + \beta\Omega_{jt} + res_{jt}. \quad (2)$$

Following Alonso-Alvarez and Molina (2023), Ω_{jt} includes the following battery of controls at different frequencies, sourced from Refinitiv. In terms of high-frequency data, we construct a volatility measure for the country-specific equity index as the 40-day standard deviation of the daily change of the domestic equity index. We also include the long-term yield of sovereign bonds, in local currency when possible, or using the Emerging Markets Bond Index (EMBI) instead. In terms of lower-frequency variables, we include the quarterly change in gross public debt as a percentage of GDP, the annual moving average of the current account balance as a percentage of GDP, countries' short-term external debt in percentage of international reserves, the inflation rate (year-on-year change of the consumer price index), and the level of central banks' international reserves in billions of USD. Finally, to proxy for the general economic performance and solvency of a country, we include the sovereign credit rating, as defined by Standard and Poor's, transforming its alphanumeric scale linearly, from 21 (AAA ratings) to 12 (BBB-, that is, the investment-grade level) and 0 (restricted defaults or selective defaults, RD and SD). Note that this approach is preferable to just controlling for country characteristics in the DID analysis on the basis since it relies on longer time series, while the DID is just

over a four-day window.¹³ Figures A.3–A.6 report the time series of the basis and the residuals.

In the last part of the paper, to assess the potential spillback effect on the EA, we consider EA banks' stock prices as a relevant metric, since they capture market valuations and expectations about bank-specific profitability. In the context of our analysis, the working hypothesis would be that the announcement of an ECB swap line with a third country outside the euro area reduces the risk of financial turbulence and/or the probability of default of banks or firms in that concrete market. This, in turn, improves the valuation of the highly exposed EA banks, and the stock price should increase. Equity prices are obtained on a daily frequency from Refinitiv, and we use the Datastream aggregate bank equity indices for each country. As in the case of the FX basis, we purge the equity indices from country-specific factors. Since in this last exercise the dependent variable is at the EA country level, we purge for variables that may affect the health of the banking sector of a given EA country. These are the short-term interest rate, defined as the three-month Treasury bill interest rate, or the closest maturity when the former is not available; the nominal effective exchange rate deviation, calculated as the difference between the observed nominal effective exchange rate and the exchange rate that would prevail if the real effective exchange rate were consistent with its long term mean, using the International Monetary Fund (IMF) nominal and real effective exchange rates; net foreign assets of domestic banks, defined as the difference between domestic banks' claims and liabilities with non-residents over GDP (as defined by the IMF's International Financial Statistics database); and the loan-to-deposit ratio, defined as domestic banks' claims on the private sector over the sum of deposits (transferable deposits included in broad money definition, other deposits included in broad money, and deposits excluded from broad money), as posted by the IMF's International Financial Statistics database. In Appendix A.1 we summarize the descriptive statistics for the stock market data as well as the banks included in the respective aggregates.

¹³Since the inclusion of low-frequency variables may create concerns about jumps in the data, as a robustness check we exclude them from the specification, and the results hold.

Finally, to construct a measure of EA countries' banking-sector exposure to foreign countries, we use the Bank for International Settlements Consolidated Banking Statistics on a quarterly basis. In particular, we consider total claims of each EA domestic bank on foreign banking sectors, all maturities, and all instruments and currencies, measured on a guarantor basis. The exposure of the banking sector of EA country j to non-EA country i is calculated as the share of claims on country i over total claims of domestic banks of country j on all countries. For example, Italy's exposure to Bulgaria is proxied by the share of Italian banks' total claims on Bulgarian banks over the total of cross-border claims of Italian banks. When missing, data are imputed using linear interpolations. The EA country exposure is then calculated as the average exposure toward the countries targeted by a liquidity facility. Based on this measure, we construct a dummy that allows us to classify EA countries' banking sectors as highly exposed to non-EA countries if the average exposure of EA country j in the quarter prior to the announcement t is higher than the 75th percentile of the cross-country and event distribution. On the contrary, we defined an EA country banking sector as lightly exposed if its exposure is lower than the 75th percentile threshold. Table A.9 in Appendix A.4 tabulates the exposure dummy across countries and announcements. As a robustness check, we also use (i) the 65th percentile as the relevant threshold and (ii) the continuous exposure measure without any threshold.

4. Direct Effect: Euro Funding in FX Markets

4.1 Mechanism

Bahaj and Reis (2022) show how an active liquidity line caps the forward price in FX swaps. Intuitively, central bank liquidity lines provide recipient-country banks with a cheaper outside option to the FX market, which improves the bargaining terms that the recipient-country banks obtain from the traders. Through the term structure, this argument holds also in the case of the mere announcement of a liquidity line, which ultimately caps the price of forward agreements.¹⁴ To see this, assume there is a trader who has

¹⁴We thank an anonymous referee for pointing this out.

access to one-day swap agreements and one-week swap agreements at time t . The liquidity announcement lowers the expected price of the future one-day agreements since there is a probability that the line will be activated, capping the swap price. Consequently, the announcement decreases also the price of one-week swaps at t .¹⁵

This section, first, presents the methodology used to estimate the signaling effect of ECB euro liquidity facilities on the euro funding cost in FX market, as measured by CIP deviations; second, it reports the results; and, finally, it provides evidence of the robustness.

4.2 Empirical Strategy

As in Bahaj and Reis (2022), our empirical strategy is based on a high-frequency DID approach. The high-frequency dimension of the approach consists of considering changes in financial variables in a short window around the announcement. It contributes to identifying the causal effect of central banks' actions since (i) markets react only to unexpected announcements and (ii) the short time window helps to exclude other confounding factors. Specifically, we consider a window spanning two days before and one day after the announcement. In addition, the DID aspect is required for two reasons. First, since these events are sporadic, it is more informative to focus on specific episodes of financial stress when the volume of trading in swap contracts increases and the equilibrium is constrained by the FX swap supply curve. Second, as underlined in Section 2 as well as in the list of events considered, the ECB euro lines have been motivated by existing FX market dysfunction. This indicates a potential intrinsic correlation with confounding global factors, such as the contemporaneous worsening of global financial conditions. To avoid biased estimates, we compare the effect in the treatment group with an adequate control group.

In our context, the selection of the comparison group for the DID analysis merits some attention. The ECB euro liquidity lines have targeted most of the currencies of central banks in the geographical vicinity of the euro. This is a standard case of multiple periods and groups, where the treatment is staggered over time

¹⁵This is proofed in Appendix C.

(Athey and Imbens 2022; Callaway and Sant’Anna 2021 among others). However, the treatment is not an absorbing state since we consider announcements and not implementations. Therefore, our identification strategy is as follows. For each announcement date, we compare treated countries, i.e., countries whose currency is targeted by the announcement of an ECB line, with a counterfactual that includes currencies that are targeted at previous or future dates.¹⁶ This strategy is similar to Fadlon and Nielsen (2021), where the authors identify the effect of health shocks in labor supply using as control group households targeted on a future date. In other words, we define a currency as treated if it is targeted by the ECB line at time t , but this same currency is considered non-treated in other periods $\hat{t} \neq t$. The idea behind this is that these countries have been targeted at least once, so they share similar characteristics under the eligibility lens of the ECB. Nevertheless, this approach could raise concerns that non-targeted countries in episode t might be affected by past liquidity line announcements. However, since we consider a short window around the announcement, we also prevent that subsequent announcements could contaminate the current one.¹⁷ Once defined the window and the control group, the identification ultimately relies on the randomness of the timing of the announcement within the window considered.¹⁸

Since we have a reduced amount of events, we collapse the panel around these announcements following previous contributions (see Bertrand, Duflo, and Mullainathan 2004), and we compute the treatment effect using the two-way fixed-effects DID estimator. More formally, we consider the following setup, in which y_t^T denotes the mean outcome of the targeted countries at time t , y_t^{NT} is the mean outcome of non-targeted countries, and t is the date of the announcement of a liquidity line. We can retrieve the treatment effect β_t by

¹⁶In the robustness we add two currencies never targeted by ECB lines.

¹⁷If liquidity line announcements had long-lasting effects that contaminate the control group and the new announcement increased the probability that agents assign to the renewal of the line in the control-group country, then our estimates would actually constitute a lower bound.

¹⁸In the robustness subsection we test for anticipation and we find no such evidence.

Figure 3. Basis Density Before and After Announcement


Note: Frequency distribution in a four-day window around the announcement. Country sample: Bulgaria, Denmark, Croatia, Hungary, Poland, Serbia, Sweden. For any announcement considered (see text in Section 3 for the list of events considered), the treated currency/ies is/are going to be the one/s targeted by the announcement, while the nontargeted currencies in the sample are untreated. Post-treatment is defined as the day of treatment and the day after, while pre-treatment is the two days prior to treatment.

comparing the outcome in t with the outcome in the previous period ($t - 1$), using the DID estimator:

$$\beta_t = (y_t^T - y_t^{NT}) - (y_{t-1}^T - y_{t-1}^{NT}). \quad (3)$$

First, as in Bahaj and Reis (2022), we inspect the distribution of the one-week FX swap basis spread around the facilities' announcements graphically. In Figure 3, we pool the observations to show the frequency distribution of the basis in this four-day window around the events for treated currencies (left panel) and for non-treated currencies (right panel). A country is considered treated in the window around event t if its currency is targeted by the ECB liquidity line announced in t , and it will be considered non-treated in the other events $\hat{t} \neq t$. As shown in Figure 3, the histogram for treated countries suggests a shift to the left. On the contrary, we do not observe

Figure 4. Residuals Before and After Announcement



Note: The figure shows the average evolution of the variable of interest in the four-day window around each announcement of an ECB liquidity line for the treated currencies versus non-treated. The variable is defined as the residual obtained from the regression of the FX basis on the recipient-country stock market volatility, change in public debt, current account balance, sovereign rating, sovereign yield, short-term debt, reserves, and inflation. For any announcement, the treated currency/ies is/are going to be the one/s corresponding to the announcement, while the rest of the currencies in the sample are used as controls. The events are pooled and the green vertical line indicates the day of the announcement.

such a shift in the non-treated group. This graphical inspection suggests that announcements of ECB liquidity lines are associated with a reduction of the cost of euro funding in the FX market.

To further reassert our grouping strategy and mimic our econometric specification, we purge the basis from country-specific relevant factors, as explained in Section 3, to control for country-specific factors that might contribute to diverging trends between control and treated currencies. Figure 4 shows the evolution of the residuals of the FX basis in the four-day window around the announcements of the treated versus the control group. The figure suggests a marked divergence at the time of the announcement ($t = 0$). Additionally, Figure A.1, in Appendix A.4, shows the evolution of the basis, which yields a similar picture.

Turning to the econometric specification, we exploit variation over two dimensions, as we do in the graphical inspection in Figure 3. The first is the time of the announcement; namely, we consider daily changes in the FX basis in the two days before the announcement versus the changes on the day of the announcement as well as the following day. Second is the variation between treatment versus the control group, i.e., currencies directly targeted by the swap line announcement relative to the rest. Equation (4) specifies this difference-in-differences (DID) framework via a two-way fixed-effect approach, where the daily change of the euro basis is regressed over a group variable for treated currencies at event e ($T_{i,e}$), which takes the value of one in the four-day interval around the announcement when the economies of concern are targeted by the ECB liquidity line, and zero for the non-targeted currencies; a period dummy ($Post_t$) that equals one on each of the days of the announcements and the following day (in t and $t + 1$), while it is zero in the two days before the announcement (in $t - 1$ and $t - 2$); as well as the interaction of the two ($T'_{i,e}Post_t$), which captures the differential effect of the line on the treated group at event e .

If the difference-in-differences estimation is implemented properly, there should be no need to control for global events that affect the treatment and control group similarly. To check for this, we include the following controls. Other ECB monetary policy decisions are captured by a dummy corresponding to the date of the ECB monetary policy meetings (mp_t). Moreover, we add a set of global controls included in a vector (Z_t), i.e., the Global Citi Economic Surprise Index and the EU high-yield spread.¹⁹

Finally, since we consider a collapsed panel, we include currency-event fixed effects $\mu_{i,e}$ to fully control for all currency-specific factors at the time of the announcement. Under this fixed-effect structure, $T_{i,e}$ is dropped because it is absorbed by the fixed effects. Standard errors are clustered at currency-event level. Finally, we include event-specific post dummies $Post_{t,e}$ to allow for a heterogeneous

¹⁹No further global volatility measures are included since the basis is already regressed on a country-specific measure of stock market volatility.

effect across events. The equation to be estimated is then given by Equation (4):

$$\begin{aligned} res_{i,t,e} = & \beta_1 T_{i,e} \times Post_t + \beta_2 T_{i,e} + \sum_e \beta_{3e} Post_{t,e} \\ & + \beta_4 mp_t + \varphi' Z_t + \mu_{i,e} + u_{i,t,e}. \end{aligned} \quad (4)$$

The effect of the ECB liquidity line is identified by β_1 , which is the group-time average treatment effect defined as the difference of the average treatment effect on the treated and control groups. The inclusion of the aforementioned global covariates controls for other events that could systematically occur in that short time window.

4.3 Results

Table 1 shows the results for the basis residuals. The euro funding cost decreases in a short window around the announcement by an average of 51 basis points in a sample of small EU countries (note that since the residuals are in percent the effect in the tables is in percentage points (pp) and we need to multiply it by 100 to interpret the effect in basis points). This effect corresponds to about 40 percent of one standard deviation of the residuals (see Table A.6 in Appendix A.4). The coefficient is invariant to controlling for monetary policy meetings and global controls. Hence, the differential evolution of CIP deviations between treatment and control groups does not appear to be explained by these variables, which strengthens the credibility of our identification strategy. The decrease in CIP deviations upon liquidity line announcements in treated countries underlines the regional importance of the ECB euro liquidity network.

Results are robust to a battery of robustness checks. First, we add Norway and Iceland to the sample. These countries can be considered comparable to some of the countries in our sample and have experienced financial distress although they have not been targeted by any ECB liquidity line. As in the baseline sample case, there is a marked divergence of the basis residuals at the time of the announcement and the following day. Empirical results are robust and reported in Table A.1, in Appendix A.2, although the magnitude of the effect reduces to 37 basis points.

Table 1. Effect of ECB Liquidity Swap Line Announcement

Dependent Variable: FX Basis Residuals		
	Baseline Sample	
	(1)	(2)
Treated $\times$ Post	−0.511** (0.030)	−0.511** (0.031)
MP Meetings		−0.211* (0.052)
EU High Yield		−0.967** (0.021)
Surprise Index		−0.00134 (0.193)
Observations	296	296
R^2	0.117	0.173
Currency $\times$ Event FE	Yes	Yes
PostMat	Yes	Yes
Note: The table reports the output of the two-way fixed-effects DID estimation on a collapsed panel which is composed by the four-day window around each ECB euro swap line announcement. Treated is defined as the currency targeted by the line in t , while the control group comprises the countries not targeted by the line announced in t . Sample composition is described in Section 4. p -values in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.		

Second, although our results are robust to the inclusion and exclusion of a battery of global controls, we further test for anticipation and potential pre-trends. We investigate whether the DID results are robust to a placebo test, in which we artificially move the announcement date three days before the actual event to detect anticipation (in case of a negative and significant effect) as well as to detect non-parallel pre-trends. By doing this, we ensure that the window in the placebo test does not include the day of the announcement. The results for the two samples can be found in Table A.3, Appendix A.2. Since the effect is always very far from being statistically significant, the results indicate no sign of anticipation and confirm the parallel pre-trends assumption.

Third, we reestimate the equation by considering the basis as dependent variable instead of the residuals. In addition to global controls, we include two country-specific covariates that change daily: stock market volatility and sovereign yield. Results are confirmed again under this setup as shown in Table A.2 in Appendix A.2.

Fourth, we check the robustness of the results using different window sizes. While augmenting the window size obliges us to drop additional events, the results remain similar (see Figure A.9 in Appendix A.4).

Fifth, we use a simple procedure to check for possible heterogeneities with respect to the intensity of treatment. This is discussed in Callaway, Goodman-Bacon, and Sant’Anna (2024), who show that heterogeneous treatment effects can be estimated by regressing outcome changes on a saturated set of “dose indicators,” where the doses in our case would correspond to the maximum volume of the liquidity line relative to (i) the size of the recipient-country central bank’s balance sheet, and (ii) international reserves of the recipient country. The latter is a natural variable to compare to the size of liquidity lines, as these, upon activation, directly increase international reserves at the disposal of the central bank to intervene in FX markets if necessary.

Let our dose measure be denoted by A_e . For the purpose of this robustness check, we estimate our baseline specification with four treatment indicators corresponding to the quartiles of A_e , $Q(A_e)$:

$$\begin{aligned}
 res_{i,t,e} = & \sum_{q=1}^4 \beta_{1,q} T_{i,e} \times \mathbf{1}(Q_{q-1}(A_e) < A_e \leq Q_q(A_e)) \\
 & \times Post_t + \sum_e \beta_{3e} Post_{t,e} + \sum_{q=1}^4 \beta_{2,q} T_{i,e} \\
 & \times \mathbf{1}(Q_{q-1}(A_e) < A_e \leq Q_q(A_e)) \\
 & + \beta_4 mp_t + \varphi' Z_t + \mu_{i,e} + u_{i,t,e},
 \end{aligned} \tag{5}$$

with $Q_0(A_e) = \min(A_e)$, $Q_4(A_e) = \max(A_e)$, and $Q_1(A_e)$, $Q_2(A_e)$, $Q_3(A_e)$ the first, second, and third quartile of A_e , respectively. The results are reported in Table A.4 in Appendix A.2, and provide some support for a role of the size of the liquidity line. The estimated treatment effect is substantially larger—approximately six to seven times greater—when the relative size of the liquidity line exceeds the third

quartile of the distribution (i.e., when the intervention is relatively large), compared to cases where it falls below the first quartile (i.e., when the intervention is relatively small).


Sixth, we estimate the treatment effect for each event separately using three different estimators: the standard two-way fixed-effects estimator we have employed so far, the Arkhangelsky et al. (2021) synthetic difference-in-differences (SDID) estimator, and the Abadie, Diamond, and Hainmueller (2015) synthetic control (SC) estimator. The latter two estimators first estimate regularized weights that optimally combine the control group observations to fit the trend of the treated units in the sample, before calculating the differencing estimator. Consistent with our previously employed strategy, treatment is defined as the day of the announcement and the day after. For each event, we keep 15 days previous to each announcement and drop any country from the control group that was treated during this time window. Figures A.7 and A.8 depict the estimated pre-trends from the Arkhangelsky et al. (2021) estimator. The average treatment effects and standard errors (in parentheses) are -0.679 (0.606) for the DID estimator, -0.259 (0.471) for the SDID estimator, and -0.484 (0.473) for the SC estimator. While the standard errors are relatively high (which is not surprising, since estimating the effect for each event separately leads to quite a low number of observations), the overall negative effect of the liquidity line announcement on the FX basis is confirmed.²⁰

Lastly, Appendix B presents alternative evidence from an event study using time-series regressions of the FX basis residuals on the post-treatment dummy and a vector of controls. The estimated effect is statistically significant and of similar magnitude as in the DID approach. The results are thus reconfirmed.

5. Spillbacks

In this section, we turn to the euro domestic market and test for spillbacks. First, we provide a theoretical argument on how the

²⁰Inference is problematic with one treated unit: We use bootstrap, since the jackknife is not defined and placebo methods outlined in Arkhangelsky et al. (2021) rely on homoskedasticity across units, which is clearly violated. We note, however, that bootstrap with just one treated unit might be less reliable. See Arkhangelsky et al. (2021) for further discussion.

Figure 5. Model Timeline (no liquidity line)

announcement of a liquidity line can increase source-country equity prices by decreasing the default probabilities of the recipient country's financial sector. Second, we test the prediction of the model empirically.

5.1 Mechanism

To illustrate how a liquidity line announcement can affect euro-area banks, we propose a stylized two-country three-period model. Each country has a representative bank. In the source country there is a euro-area bank and in the recipient country there is a local bank. The model features cross-border lending between the euro-area bank and local bank, with currency and maturity mismatches. In this context, an exogenous exchange rate shock generates a refinancing risk. The liquidity line announcement reduces the probability of default of the local bank facing the shock.²¹ Therefore, it increases the expected profits of the euro-area bank that holds claims in the recipient country. In turn, this increases its stock prices, generating positive spillbacks. Figure 5 illustrates the timeline of the model.

Endowment: Cross-Country Lending. At time $t = 0$, the euro-area bank is funded by E_0 units of equity and D_0 units of insured deposits. It holds one unit of one-period euro-denominated debt of the local banks and invested A_0 in a safe asset. The bank's balance sheet at time 0 is

$$A_0 + 1 = D_0 + E_0. \quad (6)$$

The bank is managed in the interest of equity owners, who are risk-neutral, enjoy limited liability, and have a discount factor of $\beta = 1$.

²¹The exchange rate shock is exogenous and independent of the liquidity line announcement in the model. While in reality there might be an additional positive spillback effect via the exchange rate, this is not necessary to obtain a positive spillback effect.

Shareholders receive the final asset value of the euro-area bank at $t = 2$ (if positive); no dividends are paid out before that time. We further assume that both deposits and safe assets are remunerated at a zero net interest rate, and that the bank's deposit base is constant ($D_1 = D_0$). These assumptions facilitate exposition without affecting the conclusions. The gross interest rate that the local bank has to pay to borrow from the euro-area bank is $1 + r_C$. All previous quantities are taken as exogenous from the model's perspective, as the exact quantities are irrelevant for the mechanism we wish to emphasize.²²

Meanwhile, the local bank has invested its euro-denominated debt in a euro-denominated loan maturing at time $t = 2$ with a safe net rate of return $r_L > r_C$.²³ It also has constant insured local-currency deposits D^* remunerated at net rate zero (for simplicity), which it has invested in a short-term investment I^* at a safe net return of r_I . Hence, the bank earns a safe amount $X = r_I D^*$ in local currency at $t = 1$ from this investment, implying that the recipient-country bank has funds $X + D^*$ (accrued earnings plus deposits) in local currency at $t = 1$. From the euro-area bank perspective, the $t = 0$ balance sheet and interest rates that the local bank faces are taken as exogenous. Similar to the euro-area bank, the local bank enjoys limited liability. The balance sheets of the local and the euro-area bank are illustrated in Table 2.

Exchange Rate Shock. The exchange rate S (euro per unit of local currency) is assumed to be an exogenous random variable. At $t = 1$, it is low with probability q , such that $S_L(X + D^*) < 1 + r_C$, or high with probability $1 - q$, such that $S_H(X + D^*) > 1 + r_C$. Recall that all these quantities are taken as exogenous, and neither will be affected by an unexpected announcement of a liquidity line between

²²Ex post it can be verified that there exist loan rates r_C such that it is optimal for both banks to enter the loan agreement.

²³Local banks might extend credits in source-currency denomination for various reasons, among them lower interest rates combined with underestimation of exchange rate risk (documented for, e.g., Hungarian households by Pellényi and Bilek 2009 and Indian firms by Acharya and Vij 2025), credits used to purchase or invest in foreign commodities (Caruana 2016) or portfolio allocation (Yeyati 2006).

Table 2. Model: Balance Sheets $t = 0$

Local Bank		EA Bank	
Assets	Liabilities	Assets	Liabilities
$L^* = 1$ (€, Matur. $t = 2$) I^* (LC)	1 (€, Matur. $t = 1$) D^* (LC)	A_0 (€) 1 (€)	D_0 (€) E_0 (€)

$t = 0$ and $t = 1$, the effects of which will be described below.²⁴ The exchange rate shock introduces exogenous liquidity risk for the local bank since it has to pay back the unit borrowed. If the realization of the exchange rate is low, it is unable to repay the loan to the euro-area bank using internal funds. The local bank could roll over the loan; however, to mimic the situation in which central bank liquidity lines are normally established, we assume that the market is dysfunctional and rollover is not possible.

Thus, in absence of central bank liquidity lines, the local bank only has the option to refinance the remaining debt at $t = 1$ through an FX swap. It is common knowledge that the effective cost b of refinancing one unit of debt using FX swaps is distributed according to $b \sim F_b^h$, $h \in \{nl, l\}$, where h indicates whether a liquidity line is in place (l) or not (nl).²⁵ The state dependence is discussed further below. The realization of b is learned at $t = 1$. The total cost C_s of refinancing via FX swaps is then

$$C_s(b) = (1 + b)[1 + r_c - S_L(X + D^*)] + D^*,$$

(7)

where $S_L(X + D^*)$ is the euro value of the local bank’s available funds, and D^* is its deposits. If the local bank refinances itself, it receives the repayment $1 + r_L$ of its euro-denominated loan at time $t = 2$. Otherwise, it is liquidated and receives a final payoff of zero,

²⁴In the case of the exchange rate, this is by assumption. All other quantities stated here are determined before the liquidity line is announced unexpectedly, such that they cannot be affected by the announcement.

²⁵A distribution of b obtains if as in Bahaj and Reis (2022) the local bank is randomly matched with a source-country trader and traders differ in their bargaining power. This is further explained below. While the distribution of δ_a remains fixed, a liquidity line provides an outside option to local banks, leading to the state dependence in the distribution F_b^h .

by limited liability. For simplicity, the expected exchange rate at time $t = 2$ is 1. Define then $\bar{b}$ as the maximum rate that the local bank would ever accept:

$$C_s(\bar{b}) = 1 + r_L. \quad (8)$$

Should the local bank fail to refinance the outstanding obligation with an FX swap, the euro-area bank has the right to liquidate the local bank. In this case, it receives the $t = 1$ liquidation value $\lambda_{L,1}$ euro of the long-term asset plus the bank's accrued earnings of $S_L X$ euro (deposits are owed to depositors). We assume that $\lambda_{L,1} + S_L X < 1 + r_C$, such that euro-area banks face loan default risk.

Effect on Euro-Area Bank Stock Prices without Liquidity Lines. To disentangle the effect of central bank liquidity lines on stock prices, we analyze first the counterfactual scenario without the announcement of the facility. Since equity owners are risk-neutral and do not discount the future, stock prices at the time of the liquidity line announcement, V_0 , are simply given by the expected value of time $t = 2$ equity, i.e., $V_0 = \mathbb{E}_0 V_2$. For simplicity, any funds at $t = 1$ are invested at the safe policy rate of zero, such that the euro-area bank's position in the safe asset at $t = 1$ is

$$A_1 = \begin{cases} S_L X + \lambda_{L,1} + A_0 & \text{if loan liquidated at } t = 1, \\ 1 + r_C + A_0 & \text{if loan repays at } t = 1. \end{cases} \quad (9)$$


The value of equity at time $t = 2$ is then given by

$$V_2 = \max\{A_1 - D_0, 0\}, \quad (10)$$

since the bank has to repay depositors, but enjoys limited liability. Substituting the balance sheet constraint, Equation (6), yields

$$V_2 = \max\{A_1 - 1 + E_0 - A_0, 0\}. \quad (11)$$

To compute the stock price at the time of the liquidity line announcement, it remains to characterize in which cases the euro-area bank defaults. Default of the euro-area bank can only occur at $t = 1$, since all uncertainty is realized at that time. Note that while we take r_c as exogenous, the euro-area bank would never set a loan rate r_C that

Figure 6. Model Timeline (liquidity line)

would lead it to default in case of full repayment, therefore it must be that $(1 + r_C) - 1 + E_0 \geq 0$. Furthermore, due to the possibility of default, we must have $(1 + r_C) - 1 + E_0 > 0$.

Two cases may then be distinguished: In Case 1, the euro-area bank is perfectly safe:

$$S_L X + \lambda_{L,1} - 1 + E_0 \geq 0. \quad (12)$$

Refinancing at $t = 1$ is only possible via FX swaps. Therefore, the $t = 1$ default probability of the local bank is given by

$$p_1 = q [1 - F_b^h(\bar{b})], \quad (13)$$

and thus, in Case 1 euro-area bank stock prices are given by

$$V_0 = q(1 - F_b^h(\bar{b}))(S_L X + \lambda_{L,1}) + [(1 - q) + qF_b^h(\bar{b})](1 + r_c) - 1 + E_0. \quad (14)$$

In Case 2, the euro-area bank defaults if the local bank defaults at $t = 1$, that is, if

$$S_L X + \lambda_{L,1} - 1 + E_0 < 0. \quad (15)$$

In this case, euro-area bank stock prices are given by

$$V_0 = [(1 - q) + qF_b^h(\bar{b})][(1 + r_c) - 1 + E_0]. \quad (16)$$

The euro-area bank's stock prices depend on the default probability of the local bank and, therefore, on the distribution of the CIP deviations through F_b^h .

Effect on Euro-Area Bank Stock Prices with Liquidity Lines. Now consider the announcement of a liquidity line at time $t = 0$.

Figure 6 illustrates the timeline of the model. As explained above, we have assumed that both the euro area and the local policy rates

are zero. Then, the refinancing cost via FX swaps is $b = f - s$, where f and s are the log forward rate and the log spot rate (euro per unit of local currency), respectively. Hence, b is equivalent to the FX swap basis (net of the foreign interest rate on deposits, which we have assumed to be zero for the purpose of this model). The basis, as explained in detail in the following section, captures the premium that foreign agents have to pay to borrow euros in the FX markets. Bahaj and Reis (2022) show how an active liquidity line caps the forward price in FX swaps. Intuitively, central bank liquidity lines provide local banks with a cheaper outside option to the FX market, which improves the bargaining terms that the local banks obtain from the traders. In Bahaj and Reis (2022), differences in bargaining power $\delta_a \in [0, 1]$ between source-country traders leads to a distribution of forward rates $f(\delta_a)$ for a given spot rate s , and hence a distribution of the swap basis $b(\delta_a)$, with $b(0) < b(1)$. They show that an active liquidity line reduces $b(\delta_a)$ for every bargaining power $\delta_a > 0$, i.e., unless local banks have all bargaining power. The FX basis is therefore state-dependent: $b = b^h(\delta_a)$, with $b^l(\delta_a) < b^{nl}(\delta_a) \forall \delta_a > 0$ and $b(0) = b^l(0) = b^{nl}(0)$. It follows that the swap price distribution without liquidity line has first-order stochastic dominance over the distribution with liquidity line:

$$F_b^l(x) \geq F_b^{nl}(x). \quad (17)$$

The inequality is strict for all $x \in (b(0), b^{nl}(1)]$.

As shown in Appendix C, by the theory of the term structure, this argument holds also in the case of the mere announcement of a liquidity line, which ultimately caps the price of forward agreements.

Based on the above, we can distinguish three regimes: In Regime 1, $\bar{b} < b(0)$ and the local bank defaults with probability $p_1^l = p_1^{nl} = q$, i.e., always when it receives the low exchange rate shock. Next, in Regime 2 the local bank defaults if it is matched with a trader with sufficiently high market power: $\bar{b} \in [b(0), b^{nl}(1))$ and $p_1^l, p_1^{nl} \in (0, q)$ with $p_1^l \leq p_1^{nl}$. Lastly, in Regime 3, the local bank never defaults: $\bar{b} \geq b^{nl}(1)$ and $p_1^l = p_1^{nl} = 0$.

The liquidity line announcement is therefore neutral in Regimes 1 and 3. In Regime 2, it decreases the default probability p_1 of the local bank unless it has all bargaining power.

This implies the following proposition:

PROPOSITION 1. *A liquidity line announcement in times of a freeze of the international interbank lending market leads to an increase in euro-area stock prices unless local banks exposed to a low exchange rate realization (i) never default, (ii) always default, or (iii) have all bargaining power in negotiating FX swap prices.*

5.2 Empirical Strategy

Since central bank liquidity lines are often announced in times of financial distress, the positive effect of announcements may merely cushion the negative effect of rising default probabilities due to financial distress in the recipient economy on euro funding costs and euro-area banks' balance sheets. Therefore, we might not expect to see a positive net effect of announcements in the data, but rather a positive counterfactual effect—calling for a difference-in-differences strategy. In the following exercise, we test for positive differential effects by exploiting the heterogeneous banking exposure of EA countries toward the countries considered in the treated versus control groups, in a similar fashion to Aizenman, Ito, and Pasricha (2022). While Aizenman, Ito, and Pasricha (2022) look at the impact of Fed lines on key financial variables of recipient economies, such as long-term interest rates and sovereign credit default spreads, we focus on source-country equity prices. We take advantage of the fact that some EA countries have stronger trade and banking ties with some of the countries whose currencies are targeted by ECB lines due to historical or geographical reasons. For instance, Italy has strong trade and therefore banking connections with Romania, Austria with Hungary, and so on. Therefore, we expect that an ECB announcement of a repo line with Romania will benefit Italy more than Spain, which does not have strong linkages with Romania.

While this beneficial effect may affect euro-area banks in various ways, we consider the stock price. This is an ideal measure since, by having the lowest payout priority, bank equity returns react to changes in (expected) bank profits regardless of how close a bank is to insolvency (Baron, Verner, and Xiong 2021). Consistent with our model, we expect that following an ECB liquidity line announcement the equity price of EA banks most exposed to the countries targeted by the line will increase relatively more than the less exposed banks. To account for this heterogeneous effect we consider the following

two dimensions: time (pre- versus post-announcement) and the exposure in terms of cross-border banking flows between EA source countries and non-EA counterparties. Therefore, our main explanatory variable of interest is the interaction between the dummy capturing the announcement and the following day ($Post_t$) and the exposure dummy ($Exp_{j,e}$). Furthermore, as in the case of the basis, we first purge the dependent variable for country-specific controls as explained in Section 3. Finally, we use changes and not the levels of the stock price residuals as the dependent variable, due to stationarity concerns.²⁶ We adjust the specification accordingly:

$$\begin{aligned} \Delta PriceRes_{j,t,e} = & \beta_1 Post_t \times EXP_{j,e} + \sum_e \beta_{2e} Post_{t,e} + \beta_3 EXP_{j,e} \\ & + \beta_4 mp_t + \varphi' Z_t + \mu_{j,e} + u_{j,t,e}, \end{aligned} \quad (18)$$

where $\Delta PriceRes_{j,t,e}$ is the change in the purged average log stock price of the banks in EA country j at time t in the event e , $EXP_{j,e}$ is a dummy equal to one if the average exposure of the EA country j to non-EA countries targeted by the line at event e is higher than the 75th percentile of the cross-country exposure distribution, and zero otherwise.²⁷ As in the analysis of the basis, the specification includes country-event fixed effects and a Post dummy per event. Z_t is a vector of global controls. Standard errors are clustered at country-event level.

The sample focuses on the same events of the previous exercise in order to ensure that the effect is not contaminated by liquidity lines in USD or other currencies that may happen at the same time.

Figure A.2, Appendix A.4, shows the evolution of the change in the stock price in the 10-day window around each announcement. The highly exposed and less exposed countries follow a very close pattern. Table A.9, Appendix A.4, gives an overview of the EA countries classified as exposed at each announcement. As expected,

²⁶Unit-root tests for the individual time series as well as panel unit-root tests suggest that stock prices contain a unit root. This is not surprising, as it is a common feature of financial data.

²⁷In order to simplify the notation, the exposure dummy is indexed by country and event. Precisely, it is constructed based on the country j average share in portfolio holdings by the banking sector of EA countries, in the quarter prior to the announcement t , as explained in Section 3.

Table 3. Spillbacks of Announcement to EA Banks

	(1) Banking Exposure p75	(2) Banking Exposure p65	(3) Continuous
Post × Exp	0.0210** (0.02)	0.0187** (0.04)	0.0341** (0.01)
VIX	0.000600 (0.73)	0.000633 (0.72)	0.000505 (0.77)
EU High Yield	0.0389 (0.15)	0.0390 (0.15)	0.0388 (0.15)
MP Meetings	−0.0251** (0.03)	−0.0252** (0.03)	−0.0249** (0.03)
Observations	252	252	252
R^2	0.320	0.315	0.364
Currency × Event FE	Yes	Yes	Yes
PostMat	Yes	Yes	Yes

Note: The table reports the output of the two-way fixed-effects DID estimation on a collapsed panel which is composed by the four-day window around each ECB euro liquidity line announcement. The exercise focuses on ECB lines toward EU countries to reduce potential overlapping concerns with USD liquidity lines. In columns 1 and 2 Exposure (*Exp*) is a dummy that equals one if the share of claims of the banking sector of EA-country j toward the recipient country banking sector of the country targeted by the line in t is higher than the 75th percentile (first column) or 65th percentile (second column) of the distribution across EA countries at the moment of announcement in t . In column 3, Exposure is continuous. The EA sample comprises AT, BE, DE, ES, FI, FR, IR, PT, IT. p -values in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Italy and Austria are on average more exposed than any other EA countries to ECB liquidity line recipients.

5.3 Results

Table 3 presents the results from estimating Equation (18). Consistent with the proposition of the model discussed in Section 5.1 (Proposition 1), the announcement of an ECB euro liquidity line increases stock prices of banks located in EA countries more exposed to the targeted currencies relative to banks located in less exposed

EA countries.²⁸ In the baseline specification the effect is 2.1 percent (column 1 of Table 3), corresponding to approximately 0.6 standard deviations of stock price variations. As a robustness check, instead of constructing the exposure dummy based on the 75th percentile of the distribution, we (i) reduce the threshold to the 65th percentile—Table 3, column 2—and (ii) use a continuous exposure variable—Table 3, column 3. Estimates confirm an increase in banks' stock prices. Further, the results are robust to different window sizes (Figure A.10). The model of Section 5.1 rationalizes this effect, showing that decreased euro funding costs of banks in the recipient economy due to the introduction of a liquidity line (which we provided evidence for in Section 4.3) decreases the default probability of recipient-country banks and hence increases the payoff that euro-area banks expect to receive on loans to these recipient-country banks. By improving euro-area banks' profitability, ECB liquidity lines therefore contribute to financial stability in the euro area as well. Future research at bank level may further exploit the heterogeneity in the treatment effect that ECB liquidity lines might have.

6. Conclusion

Central bank liquidity lines in times of distress function as a backstop facility, preventing episodes of liquidity shortage to turn into global financial stability problems. In line with the stated policy goal to use such swap lines as a backstop, the ECB has expanded considerably its network of swap lines during episodes of global financial stress (the Global Financial Crisis and the COVID-19 crisis).

While previous contributions have focused on the effects of USD liquidity lines, this paper presents original descriptive and empirical evidence for the case of the ECB. Concretely, we show that ECB euro liquidity lines have been effective in decreasing the premium paid by foreign agents to borrow euros in FX markets in a narrow window around the announcement. Furthermore, this paper provides evidence of positive spillbacks to the euro area generated by these

²⁸The proposition implies a positive average effect, except under the extreme assumptions that *all* banks in the targeted country either have full bargaining power in negotiating FX swap prices, or are expected to be safe with certainty, or fail with certainty in the face of adverse exchange rate shocks. Clearly, these are empirically implausible cases.

facilities in the form of relatively higher bank equity prices, which are associated with better market valuations of future profitability, in euro-area countries highly exposed via banking linkages to countries whose currencies are targeted by liquidity lines. We show how such a spillback effect on profits can arise in theory, and provide empirical evidence consistent with our theoretical predictions.

From our descriptive analysis, we notice that the ECB has provided FX insurance to central banks in its vicinity, mainly non-EA EU countries. Therefore, in contrast with the Fed, its role has been more of a regional lender of last resort than a global one. This difference seems to be driven by the different trade ties of the U.S. and the EU with the rest of the world, as supported by the evidence found in Aizenman, Ito, and Pasricha (2022) for the U.S. An unanswered question is then whether these arrangements can also boost the usage of the euro as an international currency, as has been stated in some official speeches (see Schnabel and Panetta 2020). In general, now that these lines have become a well-established tool—the ECB made its liquidity-line framework permanent in January 2024 and converted the EUREP into a standing facility open in principle to all central banks in February 2026—market participants may anticipate that liquidity in euro FX markets will be sufficient and CIP deviations may be small even in times of crisis.

Appendix A

A.1 Data Sources

Stock prices are sourced from Datastream (Refinitiv) and are country aggregates from the following banks:

- Germany: Deutsche Bank, Commerzbank, Aareal, Deutsche Pfandbriefbank, Procredit, and Umweltbank
- France: BNP, Crédit Agricole, Société Generale, Natixis, Nord CCI, Ile de France, Brie Picardie, and Crédit Foncier
- Italy: Intesa, Unicredit, Generali, BPM, BPER, Finecobank, Monte dei Paschi, Credito Emiliano, Illimity, Sondrio, Profilo, Sistema, and Piccolo credito
- Spain: Santander, BBVA, Caixabank, Bankinter, Sabadell, Liberbank, and Unicaja

- Belgium: KBC, Banque Nationale de Belgique, and KBC Ancora
- Austria: Erste, Raiffeisen, BAWAG Group, Oberbank, BKS, Addiko Bank, and Bank für Tirol und Vorarlberg
- Finland: Nordea and Aktia
- Ireland: Bank of Ireland and Permanent THB
- Portugal: Banco Comercial Portugues

A.2 Regression Tables

Table A.1. Effect of ECB Liquidity Swap Line Announcement

Dependent Variable: FX Basis Residuals				
	Baseline Sample		Extended Sample	
	(1)	(2)	(3)	(4)
Treated $\times$ Post	−0.511** (0.030)	−0.511** (0.031)	−0.373* (0.051)	−0.373* (0.052)
MP Meetings		−0.211* (0.052)		−0.0415 (0.754)
EU High Yield		−0.967** (0.021)		−0.553 (0.261)
Surprise Index		−0.00134 (0.193)		−0.00197** (0.020)
Observations	296	296	392	392
R^2	0.117	0.173	0.060	0.078
Currency $\times$ Event FE	Yes	Yes	Yes	Yes
PostMat	Yes	Yes	Yes	Yes
Note: The table reports the output of the two-way fixed-effects DID estimation on a collapsed panel which is composed by the four-day window around each ECB euro swap line announcement. Treated is defined as the currency targeted by the line in t , while the control group comprises the countries not targeted by the line announced in t . Sample composition is described in Section 4. p -values in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.				

**Table A.2. Effect of ECB Liquidity
Line Announcement (baseline sample)**

Dependent Variable: FX Basis	
	(1)
Treated $\times$ Post	−0.411** (0.026)
EU High Yield	−1.102** (0.023)
MP Meetings	−0.234** (0.027)
Surprise Index	−0.00117 (0.323)
Stock Exchange Volatility	2.465 (0.383)
Sovereign Yield	−0.838 (0.176)
Observations	304
R^2	0.204
Currency $\times$ Event FE	Yes
PostMat	Yes
Note: The table reports the output of the two-way fixed-effects DID estimation on a collapsed panel which is composed by the four-day window around each ECB euro swap line announcement. Treated is defined as the currency targeted by the line in t , while the control group comprises the countries not targeted by the line announced in t . Sample composition is described in Section 4. p -values in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.	

Table A.3. Effect of ECB Liquidity
Line Announcement (placebo)

Dependent Variable: FX Basis Residuals				
	Baseline Sample		Extended Sample	
	(1)	(2)	(3)	(4)
Treated $\times$ Post	−0.0984 (0.745)	−0.0987 (0.745)	−0.0124 (0.969)	−0.0127 (0.969)
MP Meetings		0.0629 (0.844)		0.265 (0.449)
EU High Yield		−0.346 (0.468)		−0.127 (0.759)
Surprise Index		0.00812 (0.286)		0.0103 (0.205)
Observations	296	296	392	392
R^2	0.043	0.055	0.035	0.042
Currency $\times$ Event FE	Yes	Yes	Yes	Yes
PostMat	Yes	Yes	Yes	Yes
Note: The table reports the output of the two-way fixed-effects DID estimation on a collapsed panel which is composed by the four-day window around each ECB euro swap line announcement. For the purpose of the placebo exercise, the date of the announcement has been anticipated by three days, consequently moving the window ahead of the actual announcement, preventing any overlap with true window. Treated is defined as the currency targeted by the line in t, while the control group comprises the countries not targeted by the line announced in t. Sample composition is described in Section 4. The extended sample adds Iceland and Norway to the control group. p-values in parentheses. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.				

Table A.4. Effect of ECB Liquidity Line Announcement

	Dependent Variable: FX Basis Residuals					
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	−0.511** (0.030)	−0.511** (0.031)				
Treated × Post Q1 (CB)			−0.222** (0.036)	−0.222** (0.037)		
Treated × Post Q2 (CB)			0.128 (0.148)	0.128 (0.150)		
Treated × Post Q3 (CB)			−0.719 (0.239)	−0.719 (0.241)		
Treated × Post Q4 (CB)			−1.496*** (0.005)	−1.496*** (0.005)		
Treated × Post Q1 (Reserves)					−0.233** (0.031)	−0.233** (0.032)
Treated × Post Q2 (Reserves)					0.123 (0.282)	0.123 (0.285)
Treated × Post Q3 (Reserves)					−0.961* (0.064)	−0.961* (0.065)
Treated × Post Q4 (Reserves)					−1.401** (0.049)	−1.401* (0.051)

(continued)

Table A.4. (Continued)

	Dependent Variable: FX Basis Residuals					
	(1)	(2)	(3)	(4)	(5)	(6)
MP Meetings		-0.211* (0.052)		-0.211* (0.053)		-0.211* (0.053)
EU High Yield		-0.967** (0.021)		-0.967** (0.021)		-0.967** (0.021)
Surprise Index		-0.00134 (0.193)		-0.00134 (0.195)		-0.00134 (0.195)
Observations	296	296	296	296	296	296
R^2	0.117	0.173	0.149	0.204	0.145	0.200
Currency $\times$ Event FE	Yes	Yes	Yes	Yes	Yes	Yes
PostMat	Yes	Yes	Yes	Yes	Yes	Yes

Note: The table reports the output of the two-way fixed-effects DID estimation on a collapsed panel. The dependent variable is the change in FX basis residuals with a four-day window around each ECB euro swap line announcement. The variable Treated $\times$ Post captures the effect on currencies targeted by the swap line announced at time t , while the control group comprises countries not targeted in that announcement. Sample composition is described in Section 4. The variables Treated $\times$ Post Qj, (CB) and Treated $\times$ Post Qj (Reserves), for $j \in \{1, 2, 3, 4\}$, indicate the interaction between Treated $\times$ Post and the j -th quartile of the swap line volume scaled by the size of the recipient central bank's balance sheet or by its international reserves, respectively. Quartiles are constructed by pooling the distribution of this ratio across all targeted currencies, with Q0 and Q4 corresponding to the minimum and maximum values. p -values in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.3 Timeline and Main Features of ECB Lines (Announcements)

The timeline reported in Table A.5 is based on ECB press releases. Due to space constraints, the table omits some information and links, but can be found here: [ECB liquidity lines - Timeline](#).

- The timeline does not include information on the Term Auction Facility nor the EUREP.
- Swap time extensions refer to central bank swap lines and domestic central bank repo operations against ECB-eligible collateral with domestic financial institutions to redistribute the liquidity.
- The timeline does not include discontinuations of lines or maturities.
- The column “Reciprocal” means that, in at least one of the counterparts’ press releases, there is an explicit reference to the fact that the agreement is reciprocal.
- The column “Euro Provision” indicates whether in the press release there is a clear intention of providing euro liquidity. Note that, although some agreements might be reciprocal, their actual intention is to provide liquidity in only one currency denomination.
- Publicly announced: “p” indicates lines or changes announced in ECB press releases on the same date the decision was taken, “np” lines or changes announced ex post.
- Federal Reserve (*) indicates operations and agreements in the context of the multi-central-bank agreements between the Federal Reserve, European Central Bank, Bank of England, Bank of Japan, and Swiss National Bank.
- Federal Reserve (**) indicates episodes when swap lines also enable central banks of the network of temporary bilateral liquidity swap arrangements to provide liquidity operations, should they be needed, in Japanese yen, British pound sterling, Swiss francs, Canadian dollars, and euros, in addition to U.S. dollars.
- Sveriges Riksbank (§) indicates that the arrangement was not publicly announced in 2007. The activation of the line

was announced in 2009, and the link refers to the activation announcement.

- (§§) indicates data sourced from other ECB documents different from press releases. This announcement is not included in the analysis since we assume that the main channel of communication for the signaling effect consists of press releases.
- Leg 1 reports the maximum length of drawing in the context of the first leg of the agreement, i.e., between two central banks.
- Leg 2 reports the maturities of operations with domestic counterparts, i.e., the repo facilities between recipient-country central bank to the domestic banking sector.

Table A.5. Timeline and Main Features of ECB Lines

Date	Country	Counterpart	Euro Provision	Reciprocal	Announcement	Expiration	Leg 1	Leg 2	Maximum Amount
Sep.13, 2001	USA	Federal Reserve	No	No	Swap	Oct. 13, 2001	—	—	USD 50 bn
Dec. 12, 2007	USA	Federal Reserve (*)	No	Yes	Swap	—	—	28 and 35 Days	USD 20 bn
Dec. 20, 2007	SWE	Sveriges Riksbank (§)	Yes	Yes	Swap	—	3 Months	—	EUR 10 bn
Jan. 10, 2008	USA	Federal Reserve (*)	No	Yes	Swap Time Extension	—	—	28 Days	USD 20 bn
Mar. 11, 2008	USA	Federal Reserve (*)	No	Yes	Swap Time Extension	—	—	28 Days	USD 15 bn
May 2, 2008	USA	Federal Reserve (*)	No	Yes	Swap Volume Expansion	—	—	28 Days	USD 25 bn
Sep. 18, 2008	USA	Federal Reserve (*)	No	Yes	Swap Volume Expansion	—	—	1, 28, and 84 Days	USD 40 bn
Sep. 26, 2008	USA	Federal Reserve (*)	No	Yes	Swap Volume Expansion	Over the Quarter End	—	1–7 Days	USD 35 bn
Sep. 29, 2008	USA	Federal Reserve (*)	No	Yes	Swap Volume Expansion	Apr. 30, 2009	—	—	USD 120 to 240 bn
Oct. 13, 2008	USA	Federal Reserve (*)	No	Yes	Swap Time Extension	Jan. 31, 2009	—	7, 28, and 84 Days	—
Oct. 15, 2008	CHE	Swiss National Bank	No	No	Swap	Jan. 31, 2009	—	7 Days	—
	USA	Federal Reserve (*)	No	Yes	Foreign Exchange Swaps	Jan. 31, 2009	—	—	—
Oct. 16, 2008	HUN	Magyar Nemzeti Bank	Yes	Yes	Expansion of Collateral Pool Repo	—	—	—	EUR 5 bn

(continued)

Table A.5. (Continued)

Date	Country	Counterpart	Euro Provision	Reciprocal	Announcement	Expiration	Leg 1	Leg 2	Maximum Amount
Oct. 27, 2008	DNK	Danmarks Nationalbank	Yes	Yes	Swap	—	—	—	EUR 12 bn
Nov. 11, 2008	LVA	Latvijas Banka (§§)	Yes	No	Repo	—	—	—	EUR 1 bn
Nov. 21, 2008	POL	Narodowy Bank Polski	Yes	No	Repo	—	—	—	EUR 10 bn
Dec. 19, 2008	USA	Federal Reserve (*)	No	Yes	Swap Time Extension	Mar. 31, 2009	—	7, 28, and 84 Days	—
Jan. 16, 2009	CHE	Swiss National Bank	No	No	Swap Time Extension	Apr. 30, 2009	—	7 Days	—
Feb. 3, 2009	USA	Federal Reserve (*)	No	Yes	Swap Time Extension	Oct. 30, 2009	—	—	—
Mar. 19, 2009	USA	Federal Reserve (*)	No	Yes	Swap Time Extension	Jun. 30, 2009	—	7, 28, and 84 Days	—
Apr. 6, 2009	USA	Federal Reserve (*)	Yes	Yes	Swap Volume Expansion	Oct. 30, 2009	—	—	EUR 80 bn
Jun. 10, 2009	SWE	Sveriges Riksbank	Yes	Yes	Swap Activation of 2007 Line	—	3 Months	—	EUR 3 bn (EUR 10 bn Max)
Jun. 25, 2009	USA	Federal Reserve (*)	No	Yes	Swap Time Extension	Sep. 30, 2009	—	7 and 84 Days	—
			Yes	Yes	Swap Time Extension	Feb. 1, 2010	—	—	—
	CHE	Swiss National Bank	No	No	Swap Time Extension	Oct. 31, 2009	—	7 Days	—

(continued)

Table A.5. (Continued)

Date	Country	Counterpart	Euro Provision	Reciprocal	Announcement	Expiration	Leg 1	Leg 2	Maximum Amount
Sep. 24, 2009	CHE	Swiss National Bank	No	No	Swap Time Extension	Jan. 31, 2010	—	7 Days	—
	USA	Federal Reserve (*)	No	Yes	Swap Time Extension	Jan. 31, 2010	—	7 Days	—
May 10, 2010	USA	Federal Reserve (*)	No	Yes	Swap Reactivation	—	—	7 and 84 Days	—
Dec. 17, 2010	GBR	Bank of England	No	Yes	Swap	Sep. 30, 2011	—	—	GBP 10 bn
Dec. 21, 2010	USA	Federal Reserve (*)	No	Yes	Swap Time Extension	Aug. 1, 2011	—	7 Days	—
Jun. 29, 2011	USA	Federal Reserve (*)	No	Yes	Swap Time Extension	Aug. 1, 2012	—	7 Days	—
Aug. 25, 2011	GBR	Bank of England	No	Yes	Swap Time Extension	Sep. 28, 2012	—	—	GBP 10 bn
Sep. 15, 2011	USA	Federal Reserve (*)	No	Yes	Swap Time Extension	Mar. 1, 2012	—	7 Days and 3 Months	—
Nov. 30, 2011	USA	Federal Reserve (**)	Yes	Yes	Swap Time Extension and Rate Change (Leg 2)	Feb. 1, 2013	—	7 Days and 3 Months	—
Sep. 12, 2012	GBR	Bank of England	No	Yes	Swap Time Extension	Sep. 30, 2013	—	—	—
Dec. 13, 2012	USA	Federal Reserve (**)	Yes	Yes	Swap Time Extension	Feb. 1, 2014	—	7 Days and 3 Months	—
Sep. 16, 2013	GBR	Bank of England	No	Yes	Swap Time Extension	Sep. 30, 2014	—	—	—
Oct. 10, 2013	CHN	People's Bank of China	No	Yes	Swap	Oct. 8, 2016	—	—	CNY 350 bn EUR 45 bn

(continued)

Table A.5. (Continued)

Date	Country	Counterpart	Euro Provision	Reciprocal	Announcement	Expiration	Leg 1	Leg 2	Maximum Amount
Oct. 31, 2013	USA	Federal Reserve (**)	Yes	Yes	Swap Time Extension (***)	Until Further Notice	—	7 Days and 3 Months	—
Sep. 27, 2016	CHN	People's Bank of China	No	Yes	Swap Time Extension	Oct. 25, 2019	—	—	CNY 350 bn EUR 45 bn
Mar. 5, 2019	GBR	Bank of England	Yes	Yes	Swap Activation of 2010 Line	—	—	7 Days	—
Oct. 25, 2019	CHN	People's Bank of China	No	Yes	Swap Time Extension	Oct. 8, 2022	—	—	CNY 350 bn EUR 45 bn
Mar. 15, 2020	USA	Federal Reserve (*)	No	Yes	Swap Rate Change	—	—	7 Days and 84 Days	—
Mar. 20, 2020	DNK	Danmarks Nationalbank	Yes	Yes	Swap Reactivation Swap Volume Expansion	As Long as Needed	—	—	EUR 24 bn
Apr. 15, 2020	HRV	Hrvatska Narodna Banka	Yes	No	Swap	Dec. 31, 2020	3 Months	—	EUR 2 bn
Apr. 22, 2020	BGR	Bulgarian National Bank	Yes	No	Swap	Dec. 31, 2020	3 Months	—	EUR 2 bn
Jun. 5, 2020	ROU	Banca Nationala a României	Yes	No	Repo	Dec. 31, 2020	3 Months	—	EUR 4.5 bn
Jul. 17, 2020	SRB	Narodna Banka Srbije	Yes	No	Repo	Jun. 30, 2021	3 Months	—	EUR 1 bn
	ALB	Bank of Albania	Yes	No	Repo	Jun. 30, 2021	3 Months	—	EUR 0.4 bn

(continued)

Table A.5. (Continued)

Date	Country	Counterpart	Euro Provision	Reciprocal	Announcement	Expiration	Leg 1	Leg 2	Maximum Amount
Jul. 23, 2020	HUN	Magyar Nemzeti Bank	Yes	No	Repo	Jun. 30, 2021	3 Months	—	EUR 4 bn
Aug. 18, 2020	MKD	Narodna Banka na Republika Severna Makedonija	Yes	No	Repo	Jun. 30, 2021	3 Months	—	EUR 0.4 bn
	SMR	Banca Centrale della Repubblica di San Marino	Yes	No	Repo	Jun. 30, 2021	3 Months	—	EUR 0.1 bn
Aug. 28, 2020	ROU	Banca Nationala a României	Yes	No	Repo Time Extension	Jun. 30, 2021	3 Months	—	EUR 4.5 bn
	HRV	Hrvatska Narodna Banka	Yes	No	Swap Time Extension	Jun. 30, 2021	3 Months	—	EUR 2 bn
	ALB	Bank of Albania	Yes	No	Repo Time Extension	Mar. 31, 2022	3 Months	—	EUR 0.4 bn
	HRV	Hrvatska Narodna Banka	Yes	No	Swap Time Extension	Mar. 31, 2022	3 Months	—	EUR 2 bn
	HUN	Magyar Nemzeti Bank	Yes	No	Repo Time Extension	Mar. 31, 2022	3 Months	—	EUR 4 bn

(continued)

Table A.5. (Continued)

Date	Country	Counterpart	Euro Provision	Reciprocal	Announcement	Expiration	Leg 1	Leg 2	Maximum Amount
Feb. 4, 2021	MKD	Narodna Banka na Republika Severna Makedonija	Yes	No	Repo Time Extension	Mar. 31, 2022	3 Months	—	EUR 0.4 bn
	ROU	Banca Nationala a României	Yes	No	Repo Time Extension	Mar. 31, 2022	3 Months	—	EUR 4.5 bn
	SMR	Banca Centrale della Repubblica di San Marino	Yes	No	Repo Time Extension	Mar. 31, 2022	3 Months	—	EUR 0.1 bn
	SRB	Narodna Banka Srbije	Yes	No	Repo Time Extension	Mar. 31, 2022	3 Months	—	EUR 1 bn
Mar. 28, 2022	ALB	Bank of Albania	Yes	No	Repo Time Extension	Jan. 15, 2023	3 Months	—	EUR 0.4 bn
	HUN	Magyar Nemzeti Bank	Yes	No	Repo Time Extension	Jan. 15, 2023	3 Months	—	EUR 4 bn
	MKD	Narodna Banka na Republika Severna Makedonija	Yes	No	Repo Time Extension	Jan. 15, 2023	3 Months	—	EUR 0.4 bn
	SMR	Banca Centrale della Repubblica di San Marino	Yes	No	Repo Time Extension	Jan. 15, 2023	3 Months	—	EUR 0.1 bn
Apr. 27, 2022	POL	Narodowy Bank Polski	Yes	No	Swap	Jan. 15, 2023	3 Months	—	EUR 10 bn
	ROU	Banca Nationala a României	Yes	No	Repo Time Extension	Jan. 15, 2023	3 Months	—	EUR 1 bn

Table A.7. Summary Statistics

Variable	Mean	SD	Min.	Max.	N
MP Meetings	0.038	0.191	0	1	32019
Δ Public Debt Ratio	0.16	2.48	−10.52	18.66	31608
CAB	1.541	5.852	−22.09	23.38	32019
External Debt Ratio	247.5	332.8	1.083	1839	31235
Rating	A−	4	BB−	AAA	32019
Int. Reserves	43.59	32.45	2.947	156	31749
Inflation	2.599	2.844	−2.614	18.58	31933
VIX	19.99	9.519	9.14	82.69	32019
EU High Yield	5.023	3.433	2.341	23.61	32019
Surprise Index	3.685	65.59	−304.6	212.4	32019
Sovereign Yield	3.129	2.018	−0.792	15.01	31716
Stock Exchange Volatility	1.042	0.732	0.243	11.37	32019
Note: Unbalanced panel from October 1, 2008 to September 10, 2022. CAB is the current account balance. International reserves is the level of a central bank’s international reserves in billions of USD. The external debt ratio is measured as a percentage of international reserves.					

Table A.9. Banking Exposure Dummy by EA Country and Announcement, 75th Percentile Threshold

Announcement	AT	BE	DE	ES	FI	FR	IR	IT	PT
October 16, 2008		1				0	0	1	0
October 27, 2008		0				0	1	0	1
November 21, 2008		0				1	0	1	0
June 10, 2009		0				0	1	0	1
March 20, 2020	1	0	0	0	0	0	1	0	1
April 15, 2020	1	0	0	0	0	1	0	1	0
April 22, 2020	1	1	0	0	0	1	0	0	0
July 17, 2020	1	0	0	0	0	1	0	1	0
July 23, 2010	1	1	0	0	0	0	0	1	0
August 28, 2020	1	1	0	0	0	0	0	1	0
February 4, 2021	1	1	0	0	0	0	0	1	0
March 28, 2022	1	1	1	0	0	0	0	0	0


Note: Exposure is defined as the percentage of the banking-sector claims of EA country j on non-EA country i over total claims of domestic banks of EA country j globally. If the EA country exposure is higher than or equal to the 75th percentile of the cross-country distribution at the announcement date, the exposure dummy takes a value equal to one and zero otherwise. Empty cells indicate missing values.

Table A.10. Banking Exposure Dummy by EA Country and Announcement, 65th Percentile Threshold

Announcement	AT	BE	DE	ES	FI	FR	IR	IT	PT
October 16, 2008		1				0	0	1	0
October 27, 2008		0				0	1	0	1
November 21, 2008		0				1	0	1	0
June 10, 2009		0				0	1	0	1
March 20, 2020	1	0	1	0	0	0	1	0	1
April 15, 2020	1	0	0	0	0	1	1	1	0
April 22, 2020	1	1	1	0	0	1	0	0	0
July 17, 2020	1	1	1	1	1	1	1	1	1
July 23, 1010	1	1	1	0	0	0	0	1	0
August 28, 2020	1	1	0	0	0	1	0	1	0
February 4, 2021	1	1	0	0	0	1	0	1	0
March 28, 2022	1	1	1	0	0	0	0	1	0


Note: Exposure is defined as the percentage of the banking-sector claims of EA country *j* on non-EA country *i* over total claims of domestic banks of EA country *j* globally. If the EA country exposure is higher than or equal to the 65th percentile of the cross-country distribution at the announcement date, the exposure dummy takes a value equal to one and zero otherwise. Empty cells indicate missing values.

**Figure A.1. Basis in Levels:
Averaged over Events 2008–20**




Note: The figure shows the average evolution of the FX basis in the four-day window around each announcement of an ECB liquidity line for the treated currencies versus non-treated. Each panel of the figure corresponds to a sample, as defined in Section 4. For any announcement, the treated currency/ies is/are going to be the one/s corresponding to the announcement, while the rest of the currencies in the sample are used as controls.

Figure A.2. Stock Price Changes




Note: The left panel shows the average evolution of the residuals of average bank stock prices in the 10-day window around each announcement of an ECB liquidity line for the treated currencies versus non-treated. The right panel shows the cumulative residuals since the announcement for the two groups. The residuals are obtained from regressing the average stock prices on the short-term rate, net foreign assets, loan-to-deposit ratio, and nominal effective exchange rate (NEER) deviations. Table A.9 indicates the treated and non-treated countries at each announcement considered.

Figure A.3. Time Series of CIP Deviations by Country




Note: Evolution of the FX basis over the sample period by country (Part I).

Figure A.4. Time Series of CIP Deviations by Country




Note: Evolution of the FX basis over the sample period by country (Part II).

Figure A.5. Time Series of Residuals of CIP Deviations by Country




Note: Evolution of the residuals from country-by-country regressions of the FX basis on country-specific characteristics as explained in Section 4 (Part I).

Figure A.6. Time Series of Residuals
of CIP Deviations by Country




Note: Evolution of the residuals from country-by-country regressions of the FX basis on country-specific characteristics as explained in Section 4 (Part II).

Figure A.7. Pre-trends SDID (I)



Note: The figure shows estimates of the Arkhangelsky et al. (2021) SDID estimator for all events. The dependent variable is FX Basis Residuals. Treated currencies and treatment date are indicated below each graph. Treated is defined as the currency targeted by the line at t , while the control group comprises the countries not targeted by the line announced at t , and not targeted in the previous 15 days. The specification does not include further control variables.

Figure A.8. Pre-trends SDID (II)



Note: The figure shows estimates of the Arkhangelsky et al. (2021) SDID estimator for all events. The dependent variable is FX Basis Residuals. Treated currencies and treatment date are indicated below each graph. Treated is defined as the currency targeted by the line at t , while the control group comprises the countries not targeted by the line announced at t , and not targeted in the previous 15 days. The specification does not include further control variables.

Table A.11. Summary Statistics

Country	Mean				SD				N			
	Δ Public Debt Ratio	Int. Reserves	External Debt Ratio	Rating	Δ Public Debt Ratio	Int. Reserves	External Debt Ratio	Rating (Notches)	Δ Public Debt Ratio	Int. Reserves	External Debt Ratio	Rating
BG	0.2	21.94	59.52	BBB-	1.73	6.11	28.507	1	3638	3608	3587	3638
DK	0.12	72.98	303.196	AAA	1.35	12.03	87.22	0	3638	3608	3522	3638
HR	0.68	17.61	33.081	BB+	2.42	4.21	14.136	1	3638	3608	3587	3638
HU	0.27	37.32	58.542	BBB-	2.89	8.28	9.626	1	3458	3608	3587	3638
IS	0.12	5.62	616.201	BBB+	4.71	1.45	639.21	2	3587	3608	3522	3638
NO	-0.27	62.44	405.401	AAA	2.78	10.32	71.452	0	3638	3608	3522	3638
PL	0.03	104.67	48.372	A-	1.55	20.65	10.936	0	3638	3608	3587	3638
RS	0.37	12.66	7.24	BB	1.51	1.76	4.776	1	2735	2885	2799	2915
SE	-0.02	50.85	661.435	AAA	0.81	8.21	137.678	0	3638	3608	3522	3638
All	0.48	177.87	859.395	AA-	2.58	298.37	1669.279	4	64350	64200	62280	64761


Note: Unbalanced panel from October 1, 2008 to September 10, 2022. International reserves is the level of a central bank's international reserves in billions of USD. The external debt ratio is measured as a percentage of international reserves.

Table A.12. Summary Statistics

Country	Mean				SD				N			
	CAB	Inflation	Sovereign Yield	Stock Exchange Volatility	CAB	Inflation	Sovereign Yield	Stock Exchange Volatility	CAB	Inflation	Sovereign Yield	Stock Exchange Volatility
BG	-1.08	2.57	2.37	0.86	5.324	3.55	1.42	0.56	3638	3631	3638	3638
DK	7.26	1.52	1.24	0.89	1.726	1.53	1.27	0.46	3638	3631	3638	3638
HR	-0.23	1.6	3.98	0.78	3.429	2.3	1.74	0.69	3638	3608	3335	3638
HU	0.32	3.29	4.4	1.33	2.613	2.65	1.48	0.73	3638	3631	3638	3638
IS	0.2	4.33	5.84	1.02	6.747	3.44	1.64	1.16	3638	3631	3638	3638
NO	9.22	2.34	2.2	1.27	4.568	1.2	0.95	0.8	3638	3631	3638	3638
PL	-2.16	2.76	2.9	1.09	2.503	3.02	1.16	0.54	3638	3631	3638	3638
RS	-5.85	3.93	4.15	1	2.374	3.53	1.58	0.57	2915	2908	2915	2915
SE	4.72	1.33	1.35	1.13	1.344	1.69	1.12	0.61	3638	3631	3638	3638
All	1.8	2.12	2.61	1.01	6.45	2.4	1.81	0.68	64761	64455	64458	64761


Note: Unbalanced panel from October 1, 2008 to September 10, 2022. CAB is the current account balance.

Figure A.9. Main Effect: Window Size



Note: The figure shows the coefficient β_1 as well as its standard error from estimating our baseline specification Equation (4). In the left panel, we include one day prior to the announcement in the regression, in the right panel two days. Four events are dropped from the sample to prevent overlap: April 15, 2020 and April 22, 2020; July 17, 2020 and July 23, 2020. The same events enter the sample across all window sizes.

Figure A.10. Spillback Effect: Window Size



Note: The figure shows the coefficient β_1 as well as its standard error from estimating Equation (18) using the continuous banking exposure measure (see Section 5). In the left panel, we include one day prior to the announcement in the regression, in the right panel two days. Four events are dropped from the sample to prevent overlap: April 15, 2020 and April 22, 2020; July 17, 2020 and July 23, 2020. The same events enter the sample across all window sizes.

Appendix B. Time-Series Regression

This appendix presents the results of an event study using fixed-effects estimates for the baseline and the extended country-sample of the following regression:

$$Res_{jt} = \alpha_j + \beta_1 post_{jt} + \beta \mathbf{X}_{jt} + \gamma \mathbf{W} + u_{jt}, \quad (\text{B.1})$$

where $post_{jt}$ is the main explanatory variable. It takes the value equal to one in the day of the announcement and in the following day for the treated country and zero otherwise; j indexes country, t time, $\mathbf{W}$ are monthly dummies, and $\mathbf{X}_{jt}$ is a vector of controls. This equation is estimated using fixed effects and heterogeneous robust standard errors; however, results are robust to using Driscoll-Kraay standard errors.

The results are quantitatively similar to the DID estimates presented in Section 4. Moreover, the results are robust and the estimates are slightly higher when we include week $\times$ year dummies instead of monthly fixed effects.

Table B.1. Event Study

	Baseline Sample		Extended Sample	
	(1)	(2)	(3)	(4)
Post	−0.447* (0.064)	−0.432* (0.066)	−0.484** (0.023)	−0.458** (0.020)
MP Meetings	−0.0373 (0.437)	−0.0387 (0.425)	0.00842 (0.871)	0.00538 (0.916)
EU High Yield		−0.0279 (0.507)		−0.0358 (0.297)
Surprise Index		0.00000882 (0.987)		−0.000350 (0.470)
Observations	13720	13720	17961	17961
R^2	0.045	0.046	0.074	0.074
MonthFE	Yes	Yes	Yes	Yes
CurrencyFE	Yes	Yes	Yes	Yes
Note: The table reports the output of the time-series regressions. Post takes the value of one on the day of the announcement and the day after for the currency targeted by the line in t, and zero otherwise. The extended sample adds Iceland and Norway to the control group. The model includes a fixed effect for every month in the sample, as well as currency fixed effects. p-values in parentheses. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.				

Appendix C. Proofs

C.1 Announcement Effect

The following proof shows that by the term structure, the mere announcement of ECB liquidity lines reduces the euro funding cost on the FX market.

Let $S^l(n, n+1)$ and $S^{nl}(n, n+1)$ denote the price of a swap agreement at time n maturing at time $n+1$ in the presence of a liquidity line and in absence of a liquidity line, respectively. As shown in Bahaj and Reis (2022), the price of swap agreements in presence of an active liquidity line is lower: $S^l(n, n+1) < S^{nl}(n, n+1)$. Consider now a trader choosing between entering an agreement at time 0 and maturing at time N at price $S(0, N)$, or rolling over short-term agreements. At time 0 there is no liquidity line in place, but there is a probability α_h , $h \in \{a, na\}$ that a liquidity line will be activated at any time between 0 and N , which depends on whether a line was announced ($h = a$) or not ($h = na$). Naturally, $\alpha_a > \alpha_{na}$.

The expected cost of rolling over agreements or purchasing a long-term agreement should be identical by no arbitrage, therefore:

$$S^{nl,na}(0, N) = \mathbb{E}_0 \sum_{n=1}^N (1 - \alpha_{na}) S^{nl}(n-1, n) + \alpha_{na} S^l(n-1, n) \quad (\text{C.1})$$

$$S^{nl,a}(0, N) = \mathbb{E}_0 \sum_{n=1}^N (1 - \alpha_a) S^{nl}(n-1, n) + \alpha_a S^l(n-1, n). \quad (\text{C.2})$$

It follows that

$$\begin{aligned} & S^{nl,a}(0, N) - S^{nl,na}(0, N) \\ &= \mathbb{E}_0 \sum_{n=1}^N (\alpha_a - \alpha_{na}) [S^l(n-1, n) - S^{nl}(n-1, n)]. \end{aligned} \quad (\text{C.3})$$

Hence, since $\alpha_a > \alpha_{na}$ and $S^l(n-1, n) < S^{nl}(n-1, n)$, it follows that $S^{nl,a}(0, N) < S^{nl,na}(0, N)$. ■

C.2 Proposition 1

Denote the euro-area bank's stock prices by V_0^{nl} if no liquidity line is announced, and V_0^l if a line is announced.

For Case 1, note that from Equation (14) we have

$$V_0^l - V_0^{nl} = q [F_b^l(\bar{b}) - F_b^{nl}(\bar{b})] [1 + r_c - S_L X - \lambda_{L,1}] \geq 0, \quad (\text{C.4})$$

while for Case 2, from Equation (16) we have

$$V_0^l - V_0^{nl} = q [F_b^l(\bar{b}) - F_b^{nl}(\bar{b})] [r_c + E_0] \geq 0. \quad (\text{C.5})$$

As discussed in the main text, $1 + r_c - S_L X - \lambda_{L,1} > 0$ by assumption, to ensure that euro-area banks face loan default risk. From Equation (17) we have $(F_b^l(\bar{b}) - F_b^{nl}(\bar{b})) \geq 0$. It follows that

$$F_b^l(\bar{b}) > F_b^{nl}(\bar{b}) \implies V_0^l - V_0^{nl} > 0.$$

We have explained in the main text that the minimum of $b(\delta_a)$ is attained at 0, and the maximum at 1. Further, $b(0) = b^l(0) = b^{nl}(0)$ and $b^l(1) < b^{nl}(1)$. The latter result is shown in Bahaj and Reis (2022).

By the properties of a cumulative distribution function (CDF), we then have $F_b^l(b(0)) = F_b^{nl}(b(0)) = 0$ and $F_b^l(b^l(1)) = 1$, $F_b^{nl}(b^{nl}(1)) = 1$. Therefore, $F_b^l(b^{nl}(1)) = F_b^{nl}(b^{nl}(1)) = 1$. The difference between the two CDFs can then be summarized as follows:

$$F_b^l(\bar{b}) - F_b^{nl}(\bar{b}) \begin{cases} = 0 & \text{if } \bar{b} \leq b(0), \\ > 0 & \text{if } \bar{b} \in (b(0), b^{nl}(1)), \\ = 0 & \text{if } \bar{b} \geq b^{nl}(1). \end{cases} \quad (\text{C.6})$$

Proposition 1 follows directly. ■

References

- Abadie, A., A. Diamond, and J. Hainmueller. 2015. "Comparative Politics and the Synthetic Control Method." *American Journal of Political Science* 59 (2): 495–510.
- Acharya, V. V., and S. Vij. 2021. "Regulating Carry Trades: Evidence from Foreign Currency Borrowing of Corporations in India." *Review of Economic Studies* 92 (4): 2071–2107.

- Aizenman, J., H. Ito, and G. K. Pasricha. 2022. "Central Bank Swap Arrangements in the Covid-19 Crisis." *Journal of International Money and Finance* 122 (April): Article 102555.
- Aizenman, J., Y. Jinjarak, and D. Park. 2011. "International Reserves and Swap Lines: Substitutes or Complements?" *International Review of Economics and Finance* 20 (1): 5–18.
- Aldasoro, I., T. Ehlers, P. McGuire, and G. von Peter. 2020. "Global Banks' Dollar Funding Needs and Central Bank Swap Lines." BIS Bulletin No. 27, Bank for International Settlements.
- Alonso-Alvarez, I., and L. Molina. 2023. "How To Foresee Crises? A New Synthetic Index of Vulnerabilities for Emerging Economies." *Economic Modelling* 125 (August): Article 106304.
- Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager. 2021. "Synthetic Difference-in-Differences." *American Economic Review* 111 (12): 4088–4118.
- Athey, S., and G. W. Imbens. 2022. "Design-Based Analysis in Difference-in-Differences Settings with Staggered Adoption." *Journal of Econometrics* 226 (1): 62–79.
- Baba, N., and F. Packer. 2009. "Interpreting Deviations from Covered Interest Parity During the Financial Market Turmoil of 2007–08." *Journal of Banking and Finance* 33 (11): 1953–62.
- Bahaj, S., and R. Reis. 2020. "Central Bank Swap Lines During the Covid-19 Pandemic." *Covid Economics* 2: 1–12.
- . 2022. "Central Bank Swap Lines: Evidence on the Effects of the Lender of Last Resort." *Review of Economic Studies* 89 (4): 1654–93.
- . 2023. "The Workings of Liquidity Lines Between Central Banks." In *Research Handbook of Financial Markets*, ed. R. S. Gürkaynak and J. H. Wright, 102–24. Cheltenham, UK: Edward Elgar Publishing.
- . 2026. "Jumpstarting an International Currency." *Review of Economic Studies* (advance article). <https://doi.org/10.1093/restud/rdag011>.
- Baron, M., E. Verner, and W. Xiong. 2021. "Banking Crises without Panics." *Quarterly Journal of Economics* 136 (1): 51–113.
- Bertrand, M., E. Duflo, and S. Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics* 119 (1): 249–75.

- Bordo, M. D., O. F. Humpage, and A. J. Schwartz. 2015. "The Evolution of the Federal Reserve Swap Lines since 1962." *IMF Economic Review* 63 (2): 353–72.
- Callaway, B., A. Goodman-Bacon, and P. H. Sant'Anna. 2024. "Difference-in-Differences with a Continuous Treatment." NBER Working Paper No. 32117.
- Callaway, B., and P. H. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225 (2): 200–230.
- Carré, E., and L. Le Maux. 2020. "The Federal Reserve's Dollar Swap Lines and the European Central Bank During the Global Financial Crisis of 2007–09." *Cambridge Journal of Economics* 44 (4): 723–47.
- Caruana, J. 2016. "Credit, Commodities and Currencies." Lecture at the London School of Economics and Political Science, London, February 5.
- Cetorelli, N., L. S. Goldberg, and F. Ravazzolo. 2020. "Have the Fed Swap Lines Reduced Dollar Funding Strains During the Covid-19 Outbreak?" Federal Reserve Bank of New York, *Liberty Street Economics* (blog), May 22.
- Fadlon, I., and T. H. Nielsen. 2021. "Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records." *American Economic Journal: Applied Economics* 13 (3): 1–30.
- Ferrara, G., P. Mueller, G. Viswanath-Natraj, and J. Wang. 2022. "Central Bank Swap Lines: Micro-Level Evidence." Staff Working Paper No. 977, Bank of England.
- Goldberg, L. S., C. Kennedy, and J. Miu. 2011. "Central Bank Dollar Swap Lines and Overseas Dollar Funding Costs." *Economic Policy Review* (Federal Reserve Bank of New York) 17 (1): 3–20.
- McCauley, R., and C. R. Schenk. 2020. "Central Bank Swaps Then and Now: Swaps and Dollar Liquidity in the 1960s." BIS Working Paper No. 851.
- Moessner, R., and W. A. Allen. 2010. "Central Bank Co-operation and International Liquidity in the Financial Crisis of 2008–9." BIS Working Paper No. 310.
- . 2013. "Central Bank Swap Line Effectiveness During the Euro Area Sovereign Debt Crisis." *Journal of International Money and Finance* 35 (June): 167–78.

- Obstfeld, M., J. C. Shambaugh, and A. M. Taylor. 2009. "Financial Instability, Reserves, and Central Bank Swap Lines in the Panic of 2008." *American Economic Review* 99 (2): 480–86.
- Pellényi, G., and P. Bilek. 2009. "Foreign Currency Borrowing: The Case of Hungary." FINESS Working Paper No. 5.4.
- Schnabel, I., and F. Panetta. 2020. "The Provision of Euro Liquidity Through the ECB's Swap and Repo Operations." *The ECB Blog*, August 19.
- Yeyati, E. L. 2006. "Financial Dollarization: Evaluating the Consequences." *Economic Policy* 21 (45): 62–118.

A Business Conditions Real-Time Measure of Output Gap Deviations*

Alex Ilek and Yuval Mazar
Bank of Israel

This study introduces a novel and simple method for estimating the real-time output gap using data from Israel's Business Tendency Survey. We develop an index of excess demand or supply, defined as the difference between reported supply and demand constraints. This index is incorporated into a state-space model alongside actual output to estimate the output gap. Our findings demonstrate that the derived output gap measure has a positive and statistically significant impact on both consumer price index and producer price inflation in Israel. This measure yields promising results, exhibiting explanatory power that compares favorably to several widely used alternative output gap estimation methods. Furthermore, our output gap measure demonstrates higher real-time reliability than the standard univariate method. The proposed approach complements existing methods, providing valuable insights for economic analysis and policymaking by leveraging real-time survey data and capturing businesses' perceptions of economic constraints.

JEL Codes: E31, E32, E37.

*Alex Ilek: alexei.ilek@boi.org.il, Yuval Mazar: yuval.mazar@boi.org.il. Research Department, Bank of Israel. POB 780, 91007 Jerusalem, Israel. Any views expressed in the article are those of the authors and do not necessarily reflect those of the Bank of Israel. We are grateful to the editor and two anonymous referees for their helpful comments and insights that have significantly improved the paper. We thank Amir Yaron (the Governor of the Bank of Israel) and the members of the Monetary Committee of the Bank of Israel for their useful comments. We are also grateful to Nathan Sussman, Alon Eizenberg, Sigal Ribon, Jonathan Benchimol, Eliezer Borenstein, Tim Ginker, Guy Segal, Itamar Caspi, Eyal Argov, and all other participants of the seminar at the Bank of Israel for their valuable comments and suggestions. We also thank Jonathan Vinograd for his generous assistance with the data.

1. Introduction

The measurement of the output gap is central to conducting monetary policy. The output gap represents the difference between the actual observed output and the unobserved potential output of the economy and it is a key macroeconomic indicator used by central banks to evaluate business cycles and inflationary pressures. Real-time output gap estimation is critical for effective monetary policy implementation, and designing monetary policy based on an erroneous evaluation of the output gap can result in poor monetary policy decisions and hinder the achievement of its objectives (Orphanides 2001; Orphanides and van Norden 2002; Coutiño 2016; Segal 2017).

The literature on various approaches to estimating the output gap is extensive. Canova (2025) compares the output gaps from theoretical New Keynesian dynamic stochastic general equilibrium (DSGE) models with those derived from various statistical approaches and found that the polynomial approach was the least distorting and superior to all other output gap measures considered. Álvarez and Gómez-Loscos (2018) present a comprehensive review of various univariate and multivariate estimation methods for deriving output gaps, and discuss the advantages and limitations of these methods.

The study by Błażej, Górajski, and Ulrichs (2025) is closely related to our paper. They use Statistics Poland's Business Tendency Surveys and the Annual Non-Financial Enterprises Survey to measure total economic output in the manufacturing industry by summing firm-level gross value-added. However, they assess each firm's potential production levels by applying the Hodrick-Prescott filter to the components of its production function, a method that could have significant drawbacks (Hamilton 2018).

This paper introduces a novel and simple method for estimating the real-time output gap in the economy. The approach is based on the Business Tendency Survey in Israel, which collects direct responses from firms' managers. Specifically, the survey captures these managers' subjective assessments about demand and supply constraints affecting their businesses.


The proposed output gap measure presents a complementary approach to assessing the output gap. Our key finding in this paper

demonstrates that exploiting the demand and supply constraint data from the survey significantly contributes to the output gap estimator. Additionally, we discover that this new output gap measure more effectively explains in-sample inflation in Israel compared to other methods examined. Furthermore, it exhibits greater resilience to revisions than standard univariate methods.

The intuition behind demand and supply constraints can be explained as follows: A higher supply constraint indicates that firms are experiencing a greater shortage of workers and/or raw materials and equipment needed for producing goods and services. Therefore, a higher supply constraint can be interpreted as a negative supply shift in the economy. Conversely, a higher demand constraint suggests that the demand for goods and services, both domestically and in foreign markets (exports), has weakened. Thus, a higher demand constraint can be interpreted as a negative demand shift in the economy. Consequently, when the supply constraint exceeds the demand constraint, the economy experiences excess demand. This excess demand is expected to exert upward pressure on prices. Conversely, when the demand constraint exceeds the supply constraint, the economy faces excess supply (or demand shortages), leading to downward pressure on prices. The process of translating this excess demand/supply index into an output gap measure for the economy is detailed in Section 3.3 below.

Before proceeding with our analysis, we present some basic facts about the Israeli economy. Israel is a small, open, developed economy that joined the OECD in 2008. The Israeli economy is inclined toward service production, as the weight of services in the gross domestic product is relatively high (about 70 percent) and increasing over the years. As of 2023, Israel's GDP per capita was \$53,000 in terms of purchasing power parity, which is slightly below the OECD average of \$59,000 (Figure 1). The country's population is approximately 10 million people, with an average population growth rate of 1.9 percent—among the highest in the OECD. Since 1997, Israel has experienced an average annual growth rate of 3.7 percent, with about 2.0 percent attributed to per capita growth. The Israeli economy underwent a disinflation process in the 1990s, and since the early 2000s it has operated under an inflation target regime with a range of 1–3 percent. Over the past 20 years, the average annual

Figure 1. Key Economic Indicators for Israel



inflation rate in Israel has been 1.7 percent, aligning with its inflation target range (Figure 1).

The remainder of this paper is structured as follows: Section 2 describes the data regarding demand and supply constraints, providing the foundation for our analysis. Section 3 presents our proposed method for deriving the output gap. This section includes subsection 3.4, which provides a specific description of how our approach addresses the economic impacts of COVID-19 and a military conflict that occurred toward the end of our analysis period. Section 4 offers a comprehensive comparative analysis, first comparing our resulting measure with existing methodologies, then examining the relationship between various output gap measures and both consumer and producer price inflation. This section also assesses the real-time reliability of our proposed output gap measure. Finally, Section 5 concludes.

2. Data

The Business Tendency Survey, conducted by the Central Bureau of Statistics (CBS) in Israel since 2013, provides monthly data on production constraints reported by company managers across 10 economic sectors (see Table A.1). These surveyed sectors represent approximately 85 percent of the business sector, excluding the public sector and certain smaller industries.¹ In Israel, the public sector accounts for 27 percent of total GDP and has contributed approximately 23 percent to real economic growth over the past three decades.

As an official survey by the CBS, the Business Tendency Survey adheres to international standards and employs statistical techniques to address potential response bias. Company managers rate demand and supply constraints on a four-point scale, ranging from “no constraint” to “severe constraint.”

There are two types of supply constraints: the labor supply constraint, which gauges the difficulty of recruiting employees, and the

¹The remaining 15 percent of sectors consist of smaller industries such as Electricity and Water Supply, Agriculture, Forestry, and Fishing, among others.

material supply constraint, which measures the difficulty of obtaining raw materials and equipment. The demand constraint captures the weakness of both domestic and foreign orders.

There are several key advantages of using survey data:

- **Microfounded reliability:** Company managers provide direct responses about demand and supply constraints specific to the companies they manage. This ensures that their insights are based on their intimate knowledge of their own operations, rather than on perceptions of the general state of the economy. Therefore, their responses should be considered reliable.
- **Source identification and sectoral analysis:** The survey-based measure of the output gap provides valuable insights into its underlying drivers, distinguishing between demand-side and supply-side factors. Furthermore, it enables quantification of each sector's contribution to overall excess demand or supply.
- **Timeliness and frequency:** The survey is conducted monthly, with results for the previous month reported in the first week of the current month. This allows for prompt utilization of survey data in estimating the output gap.
- **Stability of survey data and limited revision of output gap estimate:** Since managers in the survey do not revise their past responses, the index of excess demand/supply used for estimating the output gap remains unrevised. The only source of potential revision comes from updates to the GDP data.

The survey-based method, while useful, has several drawbacks. First, managers lack incentives to provide accurate responses or invest sufficient effort in giving precise answers, as there are no penalties for incorrect information (Charness, Gneezy, and Rasocha 2021). Although this could potentially bias the reported data and skew the analysis significantly, several studies indicate that using simple incentivized surveys generates only small biases relative to complex methods (Charness, Gneezy, and Rasocha 2021). Second, the limited range of response options, scaled from 0 to 3, restricts nuanced feedback. For example, if a manager considers a constraint to be “very severe,” they can only report it as a “severe constraint.” This limitation makes it impossible to distinguish between reported severe constraints and those that are truly very severe in the manager's

assessment. Third, the subjective nature of survey responses may introduce bias. Managers might disproportionately emphasize certain issues or overstate their concerns about others. To address these potential biases and the scaling issue in converting from an index to an output gap measure, we incorporate measurement errors in the survey data when estimating the output gap using a state-space model (see Section 3.3).

To provide preliminary evidence supporting the reliability of the trends survey prior to estimating the output gap, we conducted a basic regression analysis. Using quarterly GDP growth as the dependent variable and the demand constraint as an explanatory variable, we found that the demand constraint alone explains over 50 percent of GDP growth variance. The coefficient of the demand constraint is negative and highly significant, as expected. Together with global demand (proxied by import growth of OECD economies), it explains about 80 percent of the variance in quarterly GDP growth. Further regressions on four types of inflation, including demand and supply constraints, showed expected results: a positive effect for the labor supply constraint and a negative effect for the demand constraint, with similar magnitudes (see Appendix B).


3. Method

Our procedure for deriving the output gap involves three stages: First, we derive separate indices for demand and supply constraints on real economic activity, based on the Business Tendency Survey (detailed in Appendix A). Second, we define an index of excess demand/supply as the difference between the supply and demand constraints. Finally, we estimate the output gap using a state-space model that incorporates both the index of excess demand/supply from the survey and actual output data.

3.1 Stage 1: Deriving Demand and Supply Constraints

We generate separate economy-wide indices for demand and supply constraints by weighting each sector's average constraint level according to its share in business GDP (see Table A.1 in

Figure 2. Labor and Materials Supply Constraints, and a Demand Constraint: Domestic and Foreign Orders, 2013:Q1–2025:Q1



Note: The figure illustrates three key constraints derived from the Business Tendency Survey: (i) labor supply constraint (blue line), (ii) raw materials and equipment supply constraint (black line), and (iii) demand constraint (brown line).

Appendix A).² The analysis below adopts time-varying sector weights, recalculated annually from 2013 to 2024.

Figure 2 illustrates two aggregate supply constraints and the aggregate demand constraint. It is evident that the aggregate labor constraint consistently proves more severe than the constraints on

²The series of constraints used in the paper are not seasonally adjusted. The reason for this is that the Net Balance Series of the trend survey—after seasonal adjustment by the CBS—were compared to corresponding quantitative series for which the trend survey serves as a leading indicator (such as production and employment indices, and turnover indices for the trade and services sectors). This examination revealed that the correlation between the series decreased, casting doubt on the quality of the seasonal adjustment. Given that the original balances of the survey are successfully used by the Bank of Israel to forecast economic growth, the CBS decided at this stage to avoid seasonal adjustment and the publication of seasonally adjusted series.

raw materials and equipment. The average level of the aggregate labor constraint is 0.93, compared to only 0.35 for raw materials and equipment. This pattern holds true at the micro level as well: with the exception of a few outlier observations, the constraint on materials and equipment is consistently weaker than the constraint on workers across all periods and sectors examined. This result is intuitive, as labor is typically harder to acquire than materials or equipment.

3.1.1 *Constructing the Aggregate Supply Constraint*

Figure 2 illustrates two distinct supply constraints and one demand constraint. In our analysis, we aim to formally characterize the supply side of the economy by incorporating both supply constraints. To achieve this, we employ three distinct approaches. The primary differentiation among these approaches lies in the assumed degree of substitutability between labor and raw materials/equipment as production inputs.

- (i) Two factors of production are fully complementary. This can be characterized by the Leontief production function, $Y = \min \{\alpha L, \beta K\}$, where L represents labor and K represents raw materials and equipment (the parameters α and β represent fixed proportions of inputs in the production function). Given that these two inputs cannot be substituted for one another, only the most binding supply constraint should be considered—in our case, this is the labor supply constraint (as shown in Figure 2). Thus, the aggregate supply constraint, S , is determined only by the labor supply constraint, L^{const} , that is, $S = L^{const}$.
- (ii) Two factors of production are imperfect substitutes.³ In this scenario, even when the labor constraint is tighter than the equipment constraint, firms can still increase output by utilizing more equipment. Consequently, the aggregate supply constraint is determined by a combination of both the labor

³For example, this relationship can be modeled using a constant elasticity of substitution (CES) production function.

supply constraint and the constraint on raw materials and equipment: $S = \alpha L^{const} + (1 - \alpha)K^{const}$.

- (iii) Two factors of production are highly substitutable.⁴ In this scenario, when the labor supply constraint remains permanently tighter than the raw materials and equipment constraint (Figure 2), companies can exploit the high degree of substitutability by shifting toward a more capital-intensive operational structure. Consequently, the raw materials and equipment constraint becomes the primary effective constraint for the firm. Thus, the aggregate supply constraint can be expressed as $S = K^{const}$.

3.2 Stage 2: Constructing the Excess Demand/Supply Index

In the first stage, we calculate the “constraint gap” for the period from 2013:Q1 through 2025:Q1. This gap is defined as the difference between the supply constraint and the demand constraint. If the “constraint gap” is positive, then the supply constraint is more severe than the demand constraint, indicating excess demand in the economy in terms of real activity. The larger this gap, the stronger the excess demand, and vice versa.


Let us define the index of excess demand/supply as $Excess = S - D$, where S is the aggregate supply constraint and D is the aggregate demand constraint. Based on the three cases of constructing S shown in Section 3.1, we define three measures of excess demand:

- (i) $L_{excess} = L^{const} - D^{const}$, which is the gap between the labor supply constraint and the demand constraint.
- (ii) $L\&K_{excess} = \alpha L^{const} + (1 - \alpha)K^{const} - D^{const}$, which is the gap between the combined supply constraint and the demand constraint.⁵

⁴This relationship can be modeled using a linear production function.

⁵To assess the relative importance of labor versus raw materials and equipment constraints, we estimate: $\log(Y) = c + \delta \log(Trend) + \alpha \log(L^{const}) + (1 - \alpha) \log(K^{const})$, where Y is business GDP, L^{const} is the labor constraint, K^{const} is the raw materials and equipment supply constraint, and $Trend$ is a time trend. We obtain $\alpha = 0.90$, indicating that the labor constraint is significantly more important for production than the raw materials/equipment constraint. This is

Figure 3. The Evolution of Three Constraint Gap Indices, 2013:Q1–2025:Q1



Note: L_excess represents the gap between the labor supply constraint and the demand constraint. K_excess represents the gap between the raw materials and equipment supply constraint and the demand constraint. $L\&K_excess$ represents the gap between the combined supply constraint and the demand constraint.

- (iii) $K_excess = K^{const} - D^{const}$, which is the gap between the raw materials and equipment supply constraint and the demand constraint.

Figure 3 illustrates three measures of constraint gap indices. The first two indices fluctuate around zero, while the third measure, based on K_excess , remains consistently negative throughout the entire sample period. The constraint gap $L\&K_excess$ consistently falls below the constraint gap L_excess due to the incorporation of K_excess in the aggregate supply constraint. The raw materials and equipment constraint is persistently lower than the labor constraint (Figure 2), so assigning a positive weight to the former results in

higher than the 0.67 used in Argov et al. (2012) for Israel’s production function. The difference arises from our use of constraint measures rather than underlying data of production inputs. Based on these results, we calibrate $\alpha = 0.75$, balancing our findings with previous estimates and favoring a conservative approach.

a lower *LEK excess* constraint gap compared to *L excess*. In our calibration, we set a relative weight of 25 percent for the raw materials and equipment constraint. A larger weight would exert more downward pressure on the constraint gap, consequently leading to a lower estimated output gap.

Regarding the third constraint gap index based on *K excess*, it consistently remains negative throughout the observed period (see Figure 3). This persistent negativity can be primarily attributed to the consistently low values of the equipment and raw materials constraint, as illustrated in Figure 2.

In our subsequent analysis, we omit the presentation and examination of the output gap-based *K excess* for two reasons: (i) the constraint gap derived from this measure relies on an unrealistic assumption of high substitutability between labor and equipment/raw materials (as discussed in Section 3.1); and (ii) the data strongly indicate that the constraint on equipment and raw materials is not binding over the entire sample period (as shown in Figure 2). From a technical standpoint, using *K excess* would result in a negative output gap throughout the entire sample, which is not economically plausible.

3.3 Stage 3: Estimating the Output Gap

To use the constraint gap index to obtain a measure of the output gap, we construct a simple state-space model. This model consists of two main components: signal equations and state equations.

3.3.1 Signal Equations

The first signal equation (1) in our state-space model is an identity that exploits the information of the actual output to derive the output gap:

$$y_t = y_t^n + y_t^{gap}, \quad (1)$$

where y_t is the actual GDP of the business sector (in logs), y_t^n is the potential output (in logs), and y_t^{gap} is the output gap.

The second signal equation (2) exploits the information of the constraint gap index ($Excess_t$) to derive the output gap. We consider separately two types of constraint gap indices: *L excess* or

L&K excess. The unobserved output gap is linked to the chosen index with a scaling factor Υ . This relationship is expressed as

$$Excess_t = \Upsilon y_t^{gap} + \xi_t^{excess}, \quad \xi_t^{excess} \sim N(0, \sigma_{excess}^2). \quad (2)$$

The residuals ξ_t^{excess} capture the measurement errors in the constraint gap index. As noted in Section 2, the responses of survey participants might be biased or underestimated, especially during economic crises. Consequently, the precision of the signal provided by the constraint gap index is unknown a priori and depends on the variance of the residuals σ_{excess}^2 .

3.3.2 State Equations

The first state equation (3) is an identity, defining variable g as the time-variant growth rate of the potential output:

$$y_t^n = y_{t-1}^n + g_{t-1}. \quad (3)$$

The second state equation (4) describes the evolution of g as a random walk process:

$$g_t = g_{t-1} + \varepsilon_t^g, \quad \varepsilon_t^g \sim N(0, \sigma_g^2). \quad (4)$$

The last state equation (5) describes the evolution of the output gap as an $AR(1)$ process. This assumption is consistent with the data-generating process of $Excess_t$, which exhibits $AR(1)$:

$$y_t^{gap} = \beta y_{t-1}^{gap} + u_t^{gap}, \quad u_t^{gap} \sim N(0, \sigma_{ygap}^2). \quad (5)$$

It is important to note that this state-space model nests traditional measures of the output gap which rely solely on information from actual output y_t , such as the Hodrick-Prescott (HP) filter, linear and quadratic trends, and other univariate approaches.

In these traditional models, only the first signal equation ($y_t = y_t^n + y_t^{gap}$) is included, and under an appropriate choice of variance ratio, $(\frac{\sigma_{ygap}^2}{\sigma_g^2})$, one can replicate most of the univariate approaches. Under the assumption that $\Upsilon = 0$, the standard models are nested

Table 1. Estimation Results of
State-Space Model, 2013:Q1–2025:Q1

Estimated Parameter/ Constraint Gap Index	L_{excess}	$L\&K_{excess}$	Standard Model
	(1)	(2)	(3)
Υ	4.03 (0.431)	4.21 (0.393)	0
β	0.821 (0.082)	0.754 (0.099)	0.7497 (0.1364)
σ_{excess}	0.092 (0.011)	0.086 (0.0096)	0.1972 - 0.209***
σ_g	0.0011 (0.0005)	0.001 (0.0004)	0.0009 (0.0003)
σ_{ygap}	0.0284 (0.0029)	0.028 (0.0029)	0.0276 (0.003)
Log-Likelihood	128.15	132.55	98.30
Note: The first value represents the estimated parameter. The value in parentheses () represents the standard error (SE) of the estimated parameter. Columns 1–2 present the estimation results using the Kalman filter (KF) of the state-space model based on constraint gap index L_{excess} and $L\&K_{excess}$, respectively. Column 3 presents the estimation results using the KF of the state-space standard model based only on actual output (excluding constraint gap index). ***The reported values are the standard deviations of the two excess-demand indices in the data under the assumption that they contain only noise in the standard model.			

within our state-space framework. In this case, when $\Upsilon = 0$, the constraint gap index (L_{excess} or $L\&K_{excess}$) becomes pure noise ($Excess_t = \xi_t^{excess}$), unrelated to the output gap.

We estimate the model using the Kalman filter by maximum likelihood. Table 1 presents the estimation results. Column 1 shows the case where the constraint gap index is $Excess_t = L_{excess}$, while column 2 shows the case where the constraint gap index is $Excess_t = L\&K_{excess}$. Column 3 presents a nested model in which both constraint gap indices are constrained to be pure noise, which simplifies the model to a univariate standard model.

The most important parameters in our analysis are the scaling parameter, Υ , and the standard error (SE) of residuals, σ_{excess} . For both constraint gap indices, we obtain Υ values close to 4 and σ_{excess} values close to 0.09. The SE of both constraint gap indices is approximately 0.2 (as shown in column 3 of Table 1). This implies that about 60 percent of the variation of the constraint gap indices in terms of SE can be attributed to the output gap, while the remaining 40 percent is attributed to measurement errors.

A thorough understanding of the parameter Υ in Equation (2) is essential. This parameter maps the metrics of the excess demand index from the survey into the output gap. Specifically, the

output gap in the state-space model is expressed in percentage points, whereas the excess demand index has uninformed units.⁶

Given the choice of the constraint scaling, the model selects the optimal scaling parameter Υ that maximizes the likelihood function. This optimization process is constrained by the fact that the standard univariate model is nested within our expanded model and anchors the scaling of the output gap in terms of percentage points.

In our chosen scaling (between 0 and 3), an increase of one unit in the excess demand index, say from 1 to 2, would be equivalent to an increase in the output gap of 0.25 percentage points [$0.25 = 1/(\Upsilon = 4)$].⁷

To quantify the impact of incorporating the constraint gap index ($Excess_t$) into the model, we conduct a likelihood ratio test (LRT). The LRT statistic is calculated as $2 \times (LL^{OurModel} - LL^{Nested})$, where $LL^{OurModel}$ and LL^{Nested} represent the log-likelihoods of our proposed model and the nested model, respectively. Both log-likelihood values are presented in Table 1. This statistic follows a chi-square distribution with one degree of freedom, since only one restriction is imposed, $\Upsilon = 0$. We find that the LRT statistic is equal to 59.7 and 68.5, respectively, which significantly exceeds the critical value of 6.64 under 1 percent significance level. Thus, we conclude that the inclusion of the constraint gap index in the model significantly contributes to the output gap estimate.


We aim to demonstrate the contribution of the constraint gap index in two ways: First, we analyze its impact on the output gap estimation by examining (i) how it affects the smoothed estimator of the output gap, and (ii) how it influences the uncertainty of this smoothed estimator.⁸ Second, we illustrate its effectiveness in

⁶As detailed in subsection 3.2, the excess demand index is calculated as the difference between aggregate supply and demand constraints. Each constraint is arbitrarily scaled between 0 and 3, reflecting the degree of constraint severity (see Appendix A for details). The crucial point is that regardless of the chosen scaling for the constraints (e.g., 100, 200, 300, and 400 or alternatively, 0.1, 0.2, 0.3, and 0.4), a larger excess demand index indicates stronger excess demand in the economy, while a lower index indicates weaker excess demand.

⁷For alternative scaling of the excess demand index, such as 0, 100, 200, 300, the scaling parameter Υ would adjust accordingly, in this case to 400.

⁸The uncertainty of the output gap estimate is captured in this context through state uncertainty (via disturbance uncertainty).

Figure 4. Contribution of the Constraint Gap Index to the Output Gap Estimation (smoothed point estimate and standard error), 2013:Q1–2025:Q1



Note: The left-hand panel shows the difference in output gap measures between our model with the constraint gap index (Table 1, columns 1 and 2) and the model without it (Table 1, column 3). The upper line corresponds to column 1, the lower to column 2. The right-hand panel displays the ratio of standard errors (SEs) for the smoothed output gap estimate. It compares SEs from our model (Table 1, columns 1 and 2) to those from the standard model (Table 1, column 3).

explaining inflation dynamics in Section 4.2. Third, we show the real-time reliability of our output gap measure compared to a standard univariate approach in Section 4.3.

Figure 4 (left-hand panel) illustrates the difference in smoothed output gap estimates between our model and the standard model (Table 1).⁹ The figure displays two distinct lines:

- Upper line: Difference between estimates in column 1 and column 3.
- Lower line: Difference between estimates in column 2 and column 3.

Figure 4 shows that both constraint gap indices make a notable quantitative contribution to the smoothed output gap estimate, particularly at the beginning and end of the sample period. The consistently higher position of the first line compared to the second

⁹The standard model’s output gap estimate (Table 1) is similar to that from the HP filter.

is attributable to L_{excess} being consistently greater than $L\&K_{excess}$ (as illustrated in Figure 3).

Figure 4 (right-hand panel) illustrates the ratio of the SE of the smoothed output gap estimate from our model to that of the standard model (Table 1). The graph demonstrates that incorporating the constraint gap indices into our model significantly reduces the SE of the output gap estimate, with the reduction ranging from 60 percent to 70 percent.

3.4 Economic Impacts of the COVID-19 Pandemic and Recent Military Conflict


3.4.1 COVID-19 Period

It is noteworthy to examine how the supply and demand constraints evolved in response to the COVID-19 crisis, which represents one of the most significant economic disruptions in recent history. Fears of mass infection prompted governments worldwide to impose extensive lockdowns, which greatly affected global economic activity. Consequently, demand constraint increased sharply in 2020, with numerous businesses reporting significant challenges in selling their products (Figure 1).

Surprisingly, despite frequent lockdowns, company reports indicated an initial decrease in the labor supply constraint. However, this trend reversed, with the constraint beginning to increase and continuing to rise until mid-2022, marking the initial phase of economic recovery from the crisis. The equipment and raw material supply constraint showed initially limited response to the crisis. However, it subsequently increased, potentially reflecting supply chain disruptions and production challenges during the recovery period. It is worth noting that the equipment and raw material supply constraint remained considerably less pronounced than the labor supply constraint throughout this period.

The initial decline in the labor supply constraint in 2020 can be attributed to the sharp drop in demand during this period. Many companies responded by reducing or halting their recruitment efforts, consequently reporting fewer difficulties in hiring. This dynamic illustrates that when the negative demand shock hit the

Figure 5. Sectoral Contributions to the Gap between Labor Supply and Demand Constraints, 2013:Q1–2025:Q1



Note: The figure illustrates the contribution of each of the 10 sectors included in the Business Tendency Survey to the gap between the labor supply constraint and the demand constraint.

economy in 2020, the demand constraint reacted more severely than the labor supply constraint, resulting in excess supply (Figure 3).

During the COVID-19 recovery period, a notable increase in demand became evident (as indicated by a decrease in demand constraint). Simultaneously, the economy faced increasing pressure from two types of supply constraints, with the labor constraint demonstrating particular prominence (Figure 2). Consequently, the COVID-19 recovery period displayed excess demand, as measured by the difference between the labor supply constraint and the demand constraint.

Figure 5 illustrates the sectoral contributions to the difference between the labor supply constraint and the demand constraint. It highlights the notable changes in demand and supply constraints

across various sectors following the onset of COVID-19. Many sectors, particularly industry, business services, and information and communication, experienced a substantial increase in demand constraint. Interestingly, business services and industry simultaneously faced a decrease in labor constraint and were the primary contributors to excess supply during 2020. The information and communication sector, while strongly affected initially, demonstrated remarkable resilience. It was not only the first to recover but also played a pivotal role in leading the Israeli economy out of the COVID-19 crisis.

It is interesting to observe how the literature addresses the derivation of the output gap during the COVID-19 crisis. Several studies adjust their models to account for the onset of the crisis; however, these adjustments often leave the underlying factors driving the output gap measure unidentified. For example, Granados and Parra-Amado (2024) derived the output gap from a reduced-form vector autoregression (VAR) model by adding a scaling factor to residuals during the COVID-19 crisis period. Morley et al. (2023) derived the output gap from a large Bayesian VAR model, applying a similar intervention to model residuals during the COVID-19 pandemic. Júlio and Maria (2024) derived the output gap from a semistructural state-space model, in which pandemic innovations were imposed into several model equations, in addition to standard innovations. Barigozzi, Lissona, and Luciani (2024) estimated the output gap using a rich multivariate model with a common factor. To account for the economic disruptions caused by the COVID-19 pandemic, they incorporated an intervention in the common factor during this period. Their study revealed that financial variables play a crucial role in capturing the excess demand (positive output gap) during the recovery phase following the initial COVID-19 outbreak.

Our analysis, however, suggests that such interventions may not be necessary during crisis periods. This is because our method utilizes survey data, which provides real-time, up-to-date information on excess demand and its breakdown into demand and supply components. The primary limitation of this approach, however, is the potential for survey responses to be both restricted in number and subject to bias during crisis periods, as will be elaborated upon in the following section. Section 3.3 takes into account the potential bias in the survey data when estimating the output gap. Notably, despite

the methodological differences, both the aforementioned studies and our approach found a negative output gap during the COVID-19 outbreak. According to our approach, this negative output gap during 2020 is attributed to a negative demand shock.

3.4.2 Military Conflict in 2023–24

Another interesting episode occurred recently. At the end of 2023, there was a significant surge in the labor supply constraint due to the outbreak of war in Israel (Figure 2). The conflict led to massive reserve recruitment, an immediate and sharp reduction in the number of Palestinian and foreign workers, and the evacuation of many thousands of people from the northern and southern regions of Israel. During 2024, the labor supply constraint eased slightly but remained high. On the demand side, the demand constraint increased at the onset of the war in late 2023, indicating a decline in domestic demand (Figure 2). However, throughout 2024, this constraint gradually eased, suggesting a slow recovery in aggregate demand. Overall, 2024 was characterized by excess demand in the economy, primarily driven by the tight labor supply constraint. Examination of sectoral contributions in 2024 reveals that business services and construction made the most substantial positive contributions to excess demand (Figure 5). Both sectors experienced intensified supply constraints. Conversely, the commerce sector exhibited the most significant negative contribution, primarily due to an increase in the demand constraint.


4. Empirics

4.1 Evolution of Different Output Gap Estimates

Figure 6 illustrates the evolution of three estimates of the output gap in Israel from 2013:Q1 through 2025:Q1.¹⁰ The first estimate, *YGAP_L_EXCESS*, is derived from a state-space model based on the index of excess demand/supply from the survey, L_{excess} . The second estimate, *YGAP_HP*, is based on the HP filter, using a smoothing

¹⁰The sample period begins in 2013:Q1, coinciding with the initiation of the survey data collection.

Figure 6. Different Measures of the Output Gap, 2013:Q1–2025:Q1



Note: The figure presents three different output gap estimates: *YGAP_L_EXCESS* (red line, with dashed lines representing the confidence interval), estimated using a state-space model based on the index of excess demand from the survey, L_{excess} ; *YGAP_HP* (black line), derived from the HP filter; and *YGAP_PROD* (green line), derived from the production function.

parameter of 1,600. The third estimate, *YGAP_PROD*, is based on the production function and it is calculated by the Research Department of the Bank of Israel.¹¹ The output gap derived from the survey

¹¹The Cobb-Douglass production function is $Y = TFP \cdot L^\alpha K^{1-\alpha}$, where *TFP* is total factor productivity, *L* is labor input, and *K* is capital. *L* is calculated as (human capital) * (number of workers) * (average weekly hours per worker). It is calculated separately for three groups: prime-age (25–64) Israeli workers, non-prime-age (15–24 and 65+) Israeli workers, and foreign and Palestinian workers, and then aggregated to total labor input. The three groups are characterized by different levels of human capital, which is defined as $1.08^{years\ of\ education}$, where 1.08 is a (gross) return on education. *K* is calculated as the national capital stock multiplied by capital utilization *U*. The output gap is calculated as $[tfp - tfp^{HP}] + \alpha[l - l^{HP}] + (1 - \alpha)[u - \bar{u}]$, where all variables are in logs and

data shows a strong correlation with the other two measures of the output gap. However, there are notable periods of divergence.


At the start of the sample period, *YGAP_L_EXCESS* shows lower values compared to the other two measures. However, for most of the subsequent period, it exhibits higher values than *YGAP_HP* and *YGAP_PROD*. This difference is particularly pronounced during two key periods: the recovery phase following the COVID-19 crisis and the conflict in Israel in 2024. The divergent behavior can be attributed to the constraint gap, as illustrated in Figure 3. This gap plays a crucial role in the model's output gap estimate, serving as a key driving force (as detailed in Section 3.3). More importantly, during the COVID-19 recovery period and the war in Israel in 2024, the constraint gap indicates excess demand, resulting in a positive output gap. In contrast, the two alternative methods yield negative output gaps during the 2024 war, primarily reflecting the decline in actual output growth rate.

Figure 7 presents two additional statistical measures of the output gap. The first measure, *YGAP_POL*, is derived from a polynomial trend, while the second, *YGAP_CF*, employs the bandpass filter of Christiano and Fitzgerald (2003) with a 6- to 32-quarter periodicity. These measures (*YGAP_POL* and *YGAP_CF*) generally align with the estimates obtained from the HP filter and production function approaches. Moreover, they exhibit a closer resemblance to the measure derived from the state-space model based on the survey's combined supply constraint index (*YGAP_L&K_EXCESS*). However, a notable divergence is observed at the beginning of the sample period, primarily due to excess supply conditions, as illustrated in Figure 3.

Figure 8 illustrates the contrast between *YGAP_L_EXCESS* and *YGAP_L&K_EXCESS*. The disparity between these two measures is expected, as it primarily reflects the differences in constraint gaps previously observed in Figure 3. Notably, the graph reveals divergent interpretations of economic conditions in 2024. According to *YGAP_L_EXCESS*, there was significant excess demand during this period. In contrast, *YGAP_L&K_EXCESS*, while showing an

“HP” means that they are detrended by the HP filter with a smoothing factor of 1,600. The log of capital utilization (u) is demeaned by long-term average.

Figure 7. Additional Measures of the Output Gap, 2013:Q1–2025:Q1




Note: The figure displays three output gap estimates: *YGAP_L&K_EXCESS* (red line, with dashed lines indicating the confidence interval), estimated using a state-space model based on the survey’s index of excess demand, *L&K_{excess}*; *YGAP_CF* (black line), derived from the Christiano and Fitzgerald (2003) band-pass filter; and *YGAP_POL* (green line), derived from the quadratic polynomial trend.

upward trend, indicates that the output level remained close to equilibrium (near zero).

The fundamental difference between statistical univariate methods and the survey-based method lies in how their trend estimates respond to fluctuations in actual output. Statistical methods tend to have trend estimates that “smoothly follow” these fluctuations. Consequently, when actual output falls, the estimated trend also drops—albeit gradually—producing a negative output gap. Canova (2025) demonstrated the fundamental differences between the New Keynesian output gap and statistical measures, indicating that purely statistical approaches generally cannot fully replicate the New Keynesian gap.

Figure 8. Comparison of Output Gap Measures Derived from Survey Data, 2013:Q1–2025:Q1



Note: This figure compares two output gap estimates: *YGAP.L.EXCESS* (red line) and *YGAP.L&K.EXCESS* (blue line). Both are derived from state-space models using different survey indices of excess demand. *YGAP.L.EXCESS* utilizes L_{excess} , while *YGAP.L&K.EXCESS* is based on $L\&K_{excess}$.

For instance, Segal (2017) showed that in a New Keynesian model, a negative technology shock lowers potential output, but due to price rigidity, actual output falls less, resulting in a positive gap and upward inflation pressures. In contrast, statistical methods, upon observing a drop in actual output, would produce a smooth decline in the trend—leading to a negative output gap. This characteristic of pure statistical measures may lead policy-makers to draw incorrect conclusions about the output gap. Such misinterpretations could result in either overly accommodative or excessively tight monetary policy, ultimately undermining the effectiveness of monetary policy in achieving its objectives. However, a comprehensive analysis of monetary policy based on different output gap measures lies beyond the scope of this paper. This is because

changing monetary policy reactions would induce different economic outcomes, necessitating a general equilibrium model for thorough evaluation.

Although we do not directly compare our survey-based gap with the New Keynesian gap, our goal is to propose a new measure that is neither purely statistical nor fundamentally New Keynesian. The primary economic potential advantage of our approach may be attributed to its microfoundations—specifically, the responses provided by numerous company managers regarding their own firms. In this sense, it shares a microfounded perspective with the New Keynesian framework. However, unlike the New Keynesian framework, our output gap measure does not depend on a specific theoretical model.

4.2 Which Output Gap Measure Best Explains Inflation?

In this subsection, we explore the potential of our survey-based output gap measure in explaining inflation—a crucial aspect relevant to monetary policy. The theoretical relationship between inflation and the output gap (reflecting excess demand) in the New Keynesian framework has been demonstrated by studies such as Clarida, Galí, and Gertler (1999) and Galí (2015). According to this framework, inflation is determined by inflation expectations, the output gap, and cost-push shocks. Given that the relationship between prices and their determinants is derived from firms' optimization processes, it is reasonable to expect that firms' perceptions and expectations (or public expectations serving as a proxy) would be most relevant when applying the inflation equation to empirical data. In this context, our survey-based output gap measure, which reflects firm managers' perceptions of excess demand or excess supply, aligns well with these theoretical foundations.

Empirical studies demonstrate advantages in using survey-based expectations, particularly for inflation. Brissimis and Magginas (2008) found support for the forward-looking New Keynesian Phillips curve using professional forecasters' (PF) inflation expectations, which largely reduced the need for inflation indexation. Kortelainen, Paloviita, and Viren (2016) showed that PF expectations improve model performance. Similarly, Smets, Warne, and

Wouters (2014) found that incorporating PF forecasts enhances the New Keynesian model's real-time forecasting for the euro area.

We contribute to the literature by utilizing a survey-based output gap measure to explain inflation. We estimate a simple inflation equation in reduced form for the sample period 2013:Q1–2025:Q1, comparing our measure's performance with other commonly used output gap estimates. We opt for a simple, empirically oriented model rather than a New Keynesian structural model to allow greater flexibility in explaining inflation dynamics. It is crucial to note that, given our relatively short sample period, the results presented here should be interpreted as preliminary indications rather than conclusive evidence:

$$\pi_t^j = c + \sum_{i=0}^s \beta_i \Delta s_{t-i} + \sum_{i=0}^n \gamma_i \pi_{t-i}^{imp} + \sum_{i=0}^P \delta_i ygap_{t-i}^{type} + \varepsilon_t, \quad (6)$$

where $\pi_t^j = [\pi_t^a, \pi_t^b, \pi_t^c, \pi_t^p]$ and $ygap_t^{type} = [YGAP_L_EXCESS, YGAP_L\&K_EXCESS, YGAP_HP, YGAP_POL, YGAP_PROD, YGAP_CF]$. π_t^a is quarterly CPI inflation, π_t^b is quarterly CPI inflation excluding housing and fruit and vegetables, π_t^c is quarterly CPI inflation excluding fruit and vegetables and energy, and π_t^p is quarterly producer price inflation. All inflation data are seasonally adjusted. $ygap_t^{type}$ is a set of six output gap versions considered. Two first measures of the output gap are derived from the Business Tendency Survey, while four additional measures are calculated using statistical methods. We added two additional explanatory variables: the change in the exchange rate of the shekel against the dollar, Δs_t , and the change in Israel's import prices in dollars (seasonally adjusted), π_t^{imp} . These two explanatory variables are needed because Israel is a small, open economy; hence, exchange rate movements and import prices substantially impact both imported final consumption goods and imported intermediate goods.

The estimation results appear in Table 2, which reports the effects of various output gaps on inflation and highlights each gap's contribution to the regression's explanatory power (R_{adj}^2). Setting aside the survey-based output gap for a moment, the polynomial-based measure emerges as the most effective at explaining inflation, showing both the highest impact and consistent

Table 2. Estimated Effect of Various Measures of the Output Gap in Equation (6), 2013:Q1–2025:Q1

Parameter ϑ /Type of Inflation	π_t^a	π_t^b	π_t^c	π_t^p
<i>YGAP_L_EXCESS</i> ($p = 8, 10, 10, 1$)	0.06 (0.00), 7%	0.11 (0.00), 13.1%	0.03 (0.00), 9.5%	0.07 (0.04), 2.6%
<i>YGAP_L&K_EXCESS</i> ($p = 8, 10, 10, 1$)	0.06 (0.01), 6.9%	0.12 (0.00), 12.8%	0.09 (0.00), 9.0%	0.07 (0.04), 2.7%
<i>YGAP_HP</i> ($p = 6, 8, 3, 1$)	0.02 (0.04), 3.5%	0.05 (0.07), 1.2%	-0.06 (0.17), 0.8%	0.04 (0.06), 2.0%
<i>YGAP_POL</i> ($p = 8, 10, 10, 1$)	0.09 (0.01), 5.8%	0.17 (0.00), 7.7%	0.15 (0.02), 6.1%	0.04 (0.07), 2.0%
<i>YGAP_PROD</i> ($p = 6, 8, 3, 1$)	0.04 (0.08), 1.6%	0.07 (0.19), -2.1%	-0.05 (0.25), -0.6%	0.07 (0.03), 3.1%
<i>YGAP_CF</i> ($p = 4, 4, 10, 3$)	-0.05 (0.03), 4.8%	-0.05 (0.05), 3.4%	0.01 (0.02), 4.6%	-0.09 (0.13), 0.9%
R_{adj}^2 (basic)	61%	60%	43%	68%

Note: In the table, each row contains three values: (i) the estimated cumulative effect of the output gap on inflation $\sum_{i=0}^P \delta_i$ in Equation (6); (ii) the p-value of the likelihood-ratio test, which tests the redundancy of the output gap in the regression (in parentheses); and (iii) the contribution of the output gap to the explanatory power, R_{adj}^2 . In all regressions, we include the contemporaneous value and four lags of the shekel-dollar exchange rate and import prices. For each inflation type (π_t^i), we optimized the number of lags for the output gap to maximize explanatory power. The first column shows the optimal number of output gap lags for each inflation measure. The last row shows R_{adj}^2 (basic) when the output gap is omitted from Equation (6).

statistical significance. Moreover, its contribution to the regression's explanatory power is the greatest. These findings align with Canova (2025), who concludes that the polynomial approach is the least distorting and superior to the other output gap measures considered.

However, our newly introduced output gap measure (*YGAP_L_EXCESS* or *YGAP_L&K_EXCESS*) outperforms the polynomial approach in terms of boosting the regression's explanatory power, contributing to R_{adj}^2 between 2.6 percent and 13.1 percent—depending on the inflation measure used. The newly introduced output gap measure has a high contribution to CPI inflation excluding housing and fruit and vegetables and a small contribution to producer price inflation (Table 2).

Now we consider an alternative specification for the inflation equation where the dependent variable is the inflation rate over the last four quarters. This alternative specification relies on moving-average trends of the explanatory variables:

$$\begin{aligned}\pi 4_t^j = & c + \alpha \Delta s_MA(p)_t + \beta \pi_MA^{imp}(p)_t \\ & + \vartheta ygap_MA(p)_t^{type} + u_t,\end{aligned}\tag{7}$$

where $\pi 4_t^j = [\pi 4_t^a, \pi 4_t^b, \pi 4_t^c, \pi 4_t^p]$ is an average inflation rate over the last four quarters. Each explanatory variable in the regression is a moving average over the last (p) quarters, $x_MA(p)_t = \frac{1}{p} [x_t + x_{t-1} + \dots + x_{t-p}]$, where the choice of lag (p) could be different for each variable.

Table 3 presents the estimation results. The survey-based measures of the output gap show a substantially higher contribution to inflation compared to Table 2. For the first three types of inflation, the contribution is particularly significant, approaching 20 percent. However, for producer price inflation, the survey-based measure's contribution is slightly lower compared to other output gap measures. Interestingly, when examining the explanatory power of Equation (7) for producer price inflation without including the output gap, import prices and exchange rates alone account for over 90 percent of the variation. This high explanatory power inherently limits the potential contribution of any type of output gap measure from the outset.

Table 3. Estimated Effect of Various Measures of the Output Gap in Equation (7), 2013:Q1–2025:Q1

Parameter ϑ /Type of Inflation	$\pi 4_t^a$	$\pi 4_t^b$	$\pi 4_t^c$	$\pi 4_t^p$
<i>YGAP_L_EXCESS</i> ($p = 9, 10, 9, 2$)	0.09 (0.00), 22%	0.10 (0.00), 21.4%	0.08 (0.00), 20.6%	0.02 (0.41), 0.5%
<i>YGAP_L&K_EXCESS</i> ($p = 10, 9, 9, 2$)	0.10 (0.00), 22.2%	0.11 (0.00), 21.5%	0.09 (0.00), 20.7%	0.03 (0.42), 0.6%
<i>YGAP_HIP</i> ($p = 10, 10, 10, 3$)	0.15 (0.00), 10.3%	0.15 (0.00), 8.1%	0.11 (0.00), 7.5%	0.05 (0.24), 1.0%
<i>YGAP_POL</i> ($p = 10, 10, 10, 3$)	0.17 (0.00), 20.9%	0.18 (0.00), 18.4%	0.14 (0.00), 19.0%	0.06 (0.23), 1.2%
<i>YGAP_PROD</i> ($p = 10, 10, 10, 3$)	0.17 (0.00), 11.2%	0.18 (0.00), 9.3%	0.13 (0.00), 8.1%	0.07 (0.15), 1.1%
<i>YGAP_CF</i> ($p = 10, 10, 10, 1$)	0.10 (0.08), 4.2%	0.10 (0.12), 2.9%	0.06 (0.18), 1.9%	0.05 (0.07), 1.0%
R_{adj}^2 (basic)	60%	61%	54%	91%

Note: In the table, each row contains three values: (i) the estimated parameter of the output gap (ϑ) in Equation (7); (ii) the p-value of the Newey-West corrected t-test (in parentheses); and (iii) the contribution of the output gap to the explanatory power, R_{adj}^2 . For each inflation type ($\pi 4_t^i$), we optimized the number of lags for the output gap, exchange rate, and import prices to maximize R_{adj}^2 . The optimal number of lags for the exchange rate and import prices was consistently four quarters across all inflation types. The first column shows the optimal number of output gap lags for each inflation measure. The last row shows R_{adj}^2 (basic) when the output gap is omitted from Equation (7).

4.3 *Real-Time Reliability*

In this subsection, we present a brief assessment of the end-of-sample reliability of our proposed output gap estimate compared to a standard univariate model. This analysis evaluates the impact of incorporating survey data on the stability of output gap estimates in real time compared to ex post.


Our procedure involves two key steps. First, we calculate the revision in the output gap derived from the univariate standard model (Table 1, column 3) by comparing the smoothed estimate (based on full sample and final vintage GDP data) with the filtered estimate (based on real-time and first vintage GDP data). Second, we perform a similar calculation for our model, which incorporates the excess demand index from the survey (Table 1, column 1). Here, we compare the smoothed estimate (using full sample, final vintage GDP data, and survey data) with the filtered estimate (using real-time and first vintage GDP data, along with survey data).¹²

Our examination directly addresses two key sources of uncertainty in real-time output gap estimates. The first source arises from data revisions (Orphanides and van Norden 2002; Kangur et al. 2019). These revisions can significantly alter the estimated output gap over time, complicating real-time economic analysis. The second source stems from the unreliability of end-of-sample trend estimates (Orphanides and van Norden 2002; Barbarino, Berge, and Stella 2024). This issue occurs because trend-cycle decomposition methods typically require both past and future data points, which are unavailable in real time.

Several studies attempt to address the problem of unreliable end-of-sample estimates. Barbarino, Berge, and Stella (2024) suggest improving the end-sample problem by including an Okun's law relationship. Similarly, Chalmovianský and Némec (2022) demonstrate that multivariate structural approaches outperform nonstructural statistical approaches in terms of stability and assessing business cycles. However, a potential drawback of multivariate approaches is that, although the output gap derived from them is likely to

¹²To eliminate the effect of different samples in real time versus the full sample on the estimated parameters in the model, we use the model parameters from Table 1 to derive the output gaps in real time and in the full sample.

Figure 9. Comparison of Output Gap Revisions—Univariate Model vs. Survey-Based Model



Note: The blue continuous line depicts the revisions of the output gap estimate based on survey data. The black dashed line illustrates the revisions of the output gap estimate derived from a standard univariate model.

be more stable and less sensitive to end-sample problem and data revisions, these models incorporate additional unobserved variables (such as the natural rate of unemployment in an Okun’s law relationship). The uncertainty associated with the stochastic processes of these additional unobserved variables introduces more uncertainty to the output gap estimates. On a more positive note, Barigozzi and Luciani (2023) showed that using a large number of variables for the estimation of the output gap in a big sample significantly reduces the magnitude of revisions to the output gap estimates.

Returning to our analysis, Figure 9 presents a comparative visualization of the revisions between two models. The graph illustrates the revisions from the univariate model, depicted by a black dashed line, in contrast with those from our proposed model, represented by a blue continuous line. The analysis reveals that revisions in our survey-based model are more modest compared to the

univariate model. Specifically, the standard error (SE) of the revisions in our model is 1.55 percent, lower than the 1.93 percent observed in the univariate model. Furthermore, the cumulative revisions in our model amount to 4.2 percent, substantially less than the 16.4 percent in the standard model and an average revision of 0.09 percent in our model, compared to 0.33 percent in the standard model.

It is particularly noteworthy to examine the revisions during the COVID-19 period. The univariate model exhibited a significant “zigzag” pattern in its revisions throughout this period. Specifically, during the initial outbreak of the crisis in 2020, we observe substantial negative revisions ranging from 3 percent to 6 percent. Subsequently, as the recovery phase unfolded in 2022, the model showed large positive revisions of 3 percent to 4 percent.

In contrast, our model shows moderate revisions of 1.8 percent at most during the crisis outbreak in 2020, and modest revisions during the recovery period. The only exception is in the first quarter of 2021, with the highest revision of 5 percent. This overall pattern of smaller revisions suggests that our model, which incorporates survey data, provides more stable estimates of the output gap in real time, particularly during periods of economic turbulence. The enhanced stability of our model’s output gap estimates primarily stems from incorporating an additional indicator from the Business Tendency Survey. These survey data, by nature, are not subject to revisions. Moreover, they provide valuable real-time information about excess demand or supply conditions in the economy.

5. Conclusions

This study introduces a novel and simple method for deriving the real-time output gap in the economy using data from the Business Tendency Survey. Our approach complements existing methods in the literature, which generally fall into two categories: univariate methods that exploit only actual output data, and multivariate methods that often require additional unobservable variables to estimate the output gap. By incorporating survey data, our method offers a middle ground, potentially combining the simplicity of univariate approaches with the additional information typically associated with multivariate methods.

Our approach offers several key advantages: (i) easy implementation; (ii) separate estimation of demand and supply constraints as distinct sources of the output gap; and (iii) timely assessment through monthly survey data, enabling real-time evaluation of excess demand or supply. These features provide policymakers with the ability to respond more effectively and promptly to economic changes, which is crucial in uncertain environments. All this may potentially enhance the accuracy and effectiveness of monetary policy decisions.

The proposed approach can be easily implemented worldwide. Many countries already conduct Business Tendency Surveys, which could be enhanced by incorporating specific questions about demand and supply constraints faced by company managers. This addition would enable a more precise real-time assessment of excess demand or supply in the economy, providing valuable input for output gap evaluation.

Data availability: Data are available on request.

Appendix A. Calculation of Supply and Demand Constraints from the Business Tendency Survey

Here, we provide a detailed explanation of how supply and demand constraints are constructed. Each month, approximately 1,600 company managers from 10 different sectors in the business sector respond to questions specific to their companies. Among these questions, three focus on the current supply and demand constraints faced by their firms.

Labor Supply Constraint. Company managers are asked to rate the extent to which employee recruitment constraints limit their operations, using a scale of 0 to 3, where:

- 0 = no constraint
- 1 = a slight constraint
- 2 = a moderate constraint
- 3 = a severe constraint

Raw Materials and Equipment Supply Constraint. Company managers are asked to rate the extent to which raw materials

and equipment constraints limit their operations, using the same scale of 0 to 3.

Demand Constraint. Company managers are asked to rate their order demand constraints in both domestic and foreign markets (export), using the same 0 to 3 scale. For each company, an average constraint level is calculated by averaging the domestic and export order constraint ratings. Subsequently, for each sector, a simple average of these company-level averages is computed.

For each of the 10 sectors, we calculate the average constraint level by taking the simple mean of all firm-level responses within that sector. We then combine these sector-level averages into an economy-wide constraint index by constructing a weighted average, where each sector's weight corresponds to its share in the total GDP of the 10 surveyed sectors.

Two weighting schemes are employed:

- (i) time-varying shares, recalculated annually from 2013 through 2024 to capture each sector's evolving share of business GDP
- (ii) constant shares, based on the average share of business GDP for each sector over the entire 2013-2024 sample

Table A.1 presents the average sector shares for the sample period, 2013–24.

Appendix B. Effects of Demand and Labor Supply Constraints on Inflation

In this appendix, we examine the effects of demand and labor supply constraints on four types of inflation. We estimate Equation (B.1), assigning separate parameters to each constraint:

$$\pi 4_t^j = \alpha \Delta s_MA(p)_t + \beta \pi_MA^{imp}(p)_t + \tau_1 L_{MA(p)_t}^{const} + \tau_2 D_{MA(p)_t}^{const} + \varepsilon_t,$$

(B.1)

where $\pi 4_t^j = [\pi 4_t^a, \pi 4_t^b, \pi 4_t^c, \pi 4_t^p]$ represents four types of inflation, and $L_{MA(p)_t}^{const}$ and $D_{MA(p)_t}^{const}$ are moving averages of labor constraint and demand constraint, respectively.

Table B.1 presents the estimation results. For all inflation types, the supply constraint coefficient is positive (higher supply constraint represents a negative shift in supply, pushing inflation up), while the demand constraint coefficient is negative (higher demand constraint represents a negative shift in demand, pushing inflation down). Both coefficients are highly significant. The Wald test for equality of the two parameters (in absolute value) indicates that we cannot reject the null hypothesis of equal impacts at the 10 percent significance level for any of the inflation types examined.

Table B.1. Estimated Separate Effect of the Labor Supply Constraint and Demand Constraints in Equation (B.1), 2013:Q1–2025:Q1

Parameter ϑ /Type of Inflation	$\pi 4_t^a$	$\pi 4_t^b$	$\pi 4_t^c$	$\pi 4_t^p$
L_MA^{const}	1.51 (0.00)	1.82 (0.00)	1.29 (0.00)	1.15 (0.01)
D_MA^{const}	−1.31 (0.00)	−1.83 (0.00)	−1.05 (0.00)	−1.04 (0.05)
p-value of Wald Test $\tau_1 + \tau_2 = 0$	0.05	0.07	0.06	0.10

Note: For each inflation type ($\pi 4_t^i$), the table presents two values: (i) the estimated parameter of the labor supply constraint on inflation, and (ii) the estimated parameter of the demand constraint on inflation. Each parameter's p-value from the Newey-West corrected t-test appears in parentheses. The number of lags for each explanatory variable was optimally selected to maximize explanatory power ($R_{adj.}^2$). The number of lags for both labor supply and demand constraints was set equally.

References

- Álvarez, L. J., and A. Gómez-Loscos. 2018. "A Menu on Output Gap Estimation Methods." *Journal of Policy Modeling* 40 (4): 827–50.
- Argov, E., E. Barnea, A. Binyamini, E. Borenstein, D. Elkayam, and I. Rozenshtrom. 2012. "MOISE: A DSGE Model for the Israeli Economy." Bank of Israel Working Paper No. 2012.06.
- Barbarino, A., T. J. Berge, and A. Stella. 2024. "The Stability and Economic Relevance of Output Gap Estimates." *Journal of Applied Econometrics* 39 (6): 1065–81.
- Barigozzi, M., C. Lissona, and M. Luciani. 2024. "Measuring the Euro Area Output Gap." Finance and Economics Discussion Series No. 2024-099, Board of Governors of the Federal Reserve System.
- Barigozzi, M., and M. Luciani. 2023. "Measuring the Output Gap Using Large Datasets." *Review of Economics and Statistics* 105 (6): 1500–14.
- Błażej, M., M. Górajski, and M. Ulrichs. 2025. "Microdata-Based Output Gap Estimation Using Business Tendency Surveys." *Journal of Economic Dynamics and Control* 174 (May): Article 105068.
- Brissimis, S. N., and N. S. Magginas. 2008. "Inflation Forecasts and the New Keynesian Phillips Curve." *International Journal of Central Banking* 4 (2): 1–22.
- Canova, F. 2025. "FAQ: How Do I Estimate the Output Gap?" *Economic Journal* 135 (665): 59–80.
- Chalnovianský, J., and D. Němec. 2022. "Assessing Uncertainty of Output Gap Estimates: Evidence from Visegrad Countries." *Economic Modelling* 116: Article 105994.
- Charness, G., U. Gneezy, and V. Rasocha. 2021. "Experimental Methods: Eliciting Beliefs." *Journal of Economic Behavior & Organization* 189: 234–56.
- Christiano, L. J., and T. J. Fitzgerald. 2003. "The Band Pass Filter." *International Economic Review* 44 (2): 435–65.
- Clarida, R., J. Galí, and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (4): 1661–707.

- Coutiño, A. 2016. "Pitfalls in Monetary Policy Decisions Based on the Output Gap." *Journal of Policy Modeling* 38 (1): 54–64.
- Galí, J. 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*, 2nd ed. Princeton University Press.
- Granados, C., and D. Parra-Amado. 2024. "Estimating the Output Gap after COVID: How to Address Unprecedented Macroeconomic Variations." *Economic Modelling* 135 (C): Article 106711.
- Hamilton, J. D. 2018. "Why You Should Never Use the Hodrick-Prescott Filter." *Review of Economics and Statistics* 100 (5): 831–43.
- Júlio, P., and J. R. Maria. 2024. "Trends and Cycles during the COVID-19 Pandemic Period." *Economic Modelling* 139 (C): Article 106830.
- Kangur, A., K. Kirabaeva, J.-M. Natal, and S. Voigts. 2019. "How Informative Are Real Time Output Gap Estimates in Europe?" IMF Working Paper No. 2019/200.
- Kortelainen, M., M. Paloviita, and M. Viren. 2016. "How Useful Are Measured Expectations in Estimation and Simulation of a Conventional Small New Keynesian Macro Model?" *Economic Modelling* 52 (B): 540–50.
- Morley, J., D. Rodríguez-Palenzuela, Y. Sun, and B. Wong. 2023. "Estimating the Euro Area Output Gap Using Multivariate Information and Addressing the COVID-19 Pandemic." *European Economic Review* 153 (C): Article 104385.
- Orphanides, A. 2001. "Monetary Policy Rules Based on Real-Time Data." *American Economic Review* 91 (4): 964–85.
- Orphanides, A., and S. van Norden. 2002. "The Unreliability of Output-Gap Estimates in Real Time." *Review of Economics and Statistics* 84 (4): 569–83.
- Segal, G. 2017. "To Respond or Not to Respond: Measures of the Output Gap in Theory and in Practice." *International Journal of Central Banking* 13 (2): 73–120.
- Smets, F., A. Warne, and R. Wouters. 2014. "Professional Forecasters and Real-Time Forecasting with a DSGE Model." *International Journal of Forecasting* 30 (4): 981–95.

“El Niño” and “La Niña”: Revisiting the Impact on Food Commodity Prices and Euro-Area Consumer Prices*

Fructuoso Borrallo, Lucía Cuadro-Sáez,
Corinna Ghirelli, and Javier J. Pérez
Banco de España

This paper challenges the assumption that the intensification of El Niño and La Niña generally exert upward and downward pressures, respectively, on international food commodity prices that, in turn, affect consumer prices even in Europe. We show that the impact is weaker nowadays and, in some cases, may even change sign. Indeed, our results suggest that El Niño may not have contributed to the upward inflation pressures in the euro area during the most recent inflationary episode. Considering domestic agricultural policies is key for the results. We construct two databases: (i) a comprehensive granular database covering both euro-area consumer prices and European Union food commodity prices; (ii) a new database of EU prices from 1972 onward, which allows us to test the effect of these events on a larger sample. We assess the impact of these climatic events by means of local projections techniques. Our bottom-up approach, which factors in disaggregated effects on

*We thank Oscar Bajo, Diego Rodríguez, Carlos Llano, Paola González, Antonio Myro, Pedro Mejía, Juan de Lucio, Rafael Myro, Gonzalo Solana, Florens Odendahl, Rodolfo Campos, and Matias Pacce for their useful comments. We also thank participants at a Banco de España internal seminar and the seminar participants at the NCB4 meeting (Frankfurt), the VI Encuentro del Banco de España con la Red de Investigación en Internacionalización (Madrid), the European System of Central Banks Working Group on Forecasting meeting, the 2024 CEMLA conference (Mexico), and a European Commission – Joint Research Center seminar for helpful discussions and comments. We also thank José Camas, Eduardo Alfaro, and Alba Villahermosa for their excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily represent the views of Banco de España or the Eurosystem. Author emails: fructuoso.borrallo@bde.es, lucia.cuadro@bde.es, corinna.ghirelli@bde.es, and javierperez@bde.es.

specific food components to assess the overall effect of El Niño and La Niña on aggregated prices, provides a more nuanced understanding of these phenomena. Furthermore, armed with our databases, we can replicate most of the existing evidence on international food commodity prices, thereby validating our approach and indirectly testing the robustness of our main results based on the euro area.

JEL Codes: C32, F62, F64, O13, Q54.

1. Introduction

The 2023–24 El Niño peaked as one of the strongest events on record, comparable to the extraordinary El Niño episode of 1997–98, according to the World Meteorological Organization (WMO) and in line with predictions by the National Center for Atmospheric Research (Yeager et al. 2022).¹ El Niño is a cyclically occurring climate phenomenon associated with warming of the ocean surface in the central and eastern tropical Pacific Ocean.² It induces anomalous temperatures, precipitation, and speeds and strength of ocean currents, leading to extended episodes of drought and floods in tropical countries.³


The counterpart phenomenon to El Niño, La Niña, induces the opposite effect, i.e., cooler sea surface temperatures across the central and eastern equatorial Pacific. La Niña sometimes follows El Niño events. Together, El Niño and La Niña represent the “warm” (El Niño) and “cold” (La Niña) phases of the El Niño–Southern Oscillation (ENSO). These climate alterations significantly affect agricultural yields in affected regions (e.g., Iizumi et al. 2014, Hsiang and Meng 2015, Smith and Ubilava 2017). In addition, indirect effects on the economy may stem from trade and relate to global value chain production and food commodity market competitiveness. Thus, ENSO has recently attracted significant attention from policymakers and central bankers because of its possible adverse (upward) effects on prices. In particular, in the last two years this

¹The National Center for Atmospheric Research (NCAR) is supported by the National Science Foundation. More information can be found here.

²It occurs every two to seven years on average and typically lasts 9 to 12 months.

³See Zebiak et al. (2015) and the literature therein.

Figure 1. Newspaper-Based ENSO Popularity Indices



Note: The figure shows newspaper-based ENSO popularity indicators for specific countries at a yearly frequency. These indicators are constructed following the procedure described in Appendix M. Red-shaded (blue-shaded) areas indicate strong El Niño (La Niña) episodes, which are declared when average sea surface temperatures in the equatorial Pacific are at least 1.5° Celsius warmer (cooler) than normal. Strong El Niño episodes occurred in 1997–98, 2009–10, 2014–16, and 2023. Strong La Niña episodes occurred in 1998–2001, 2007–08, and 2010–11. For more details, see Table 2 in Section 7.

phenomenon has been in the spotlight in the press as never before. Figure 1 shows newspaper-based indicators that capture the popularity of ENSO phenomena in the press in several developed countries.⁴ As expected, the press draws attention to ENSO events when they occur. However, the tremendous peak corresponding to the last El Niño event (2023–24) stands out. This confirms that the public is now extremely interested in this topic.

This paper investigates the effect of El Niño and La Niña events on euro-area inflation. This issue is particularly of interest today since most economies are getting back on track after the 2022–23 inflationary crisis. In this context, we find it extremely important to properly present and discuss the implications of ENSO events for the euro area. For instance, the European Central Bank (ECB) has

⁴The United States, Spain, Italy, Germany, and France. Note that these European countries are not directly affected by ENSO events, yet they talk about them when they occur.

regularly mentioned El Niño as a risk factor in its inflation forecasts. To give an example, in July 2023 the ECB stated: “It was observed that, while food inflation had started to come down somewhat, it remained elevated overall, and geopolitical risks and weather and climate-related factors, including El Niño, pointed to upside risks.”⁵ Moreover, in its December 2023 macroeconomic projections, the ECB warns about the “combination of climate change-related risks and the El Niño phenomenon which might amplify the frequency of extreme weather events resulting from climate change in general.” See Box 4 in ECB (2023).

This conjecture is supported by rich evidence on the impact of El Niño events, both from a global perspective (e.g., Brunner 2002; Hsiang and Meng 2015; Cashin, Mohaddes, and Raissi 2017; Adolfson and Lappe 2023) and also based on developing countries, which are directly affected by these phenomena (e.g., Martín 2016; Abril-Salcedo, Melo-Velandia, and Parra-Amado 2020). This evidence supports the broad consensus that El Niño raises food commodity prices, although country-specific consumer price impacts may differ depending on a country’s exposure to these climate patterns, its openness to global commodity trade, and its food commodity trade orientation (e.g., Cashin, Mohaddes, and Raissi 2017). These results are reported in detail in Table C.2 in Appendix C. Thus, there are differences in the impact of El Niño events among countries (e.g., Cashin, Mohaddes, and Raissi 2017; Andrian et al. 2024) and among food items (e.g., Ubilava 2011, 2017; Ubilava and Holt 2013; Iizumi et al. 2014; Smith and Ubilava 2017) that should be taken into account. However, to our knowledge, the literature has not yet studied the impact of ENSO events on euro-area consumer prices. Only a few studies focus on prices in specific European countries or some form of European aggregate, but not the euro area.⁶ This paper aims to fill this gap.

⁵See “Account of the Monetary Policy Meeting of the Governing Council of the ECB of 14–15 June 2023.”

⁶Brunner (2002) focuses on the G7 aggregate; Cashin, Mohaddes, and Raissi (2017) consider an ad hoc definition of the European region that includes the following 13 countries: Austria, Belgium, Finland, France, Germany, Italy, Netherlands, Norway, Spain, Sweden, Switzerland, Turkey, and the United Kingdom; Laosuthi and Selover (2007) study a sample of developed countries that includes France, Italy, and Germany.

Our main finding is twofold. First, we find that El Niño shocks increased commodity prices in the 1980s, in line with previous aforementioned results, while La Niña entailed a weaker inflationary effect on food commodity prices.⁷ Second, from 1997 onward, ENSO shocks have led to asymmetric effects on euro-area inflation. Specifically, La Niña exerts an inflationary effect both on (international and EU) food commodity prices and euro-area inflation. In contrast, El Niño shocks negatively affect EU food commodity prices, and this effect passes through (partially) to euro-area consumer prices. Hence, El Niño shocks have disinflationary effects when considering the most recent sample. These results are novel given that the literature on ENSO shocks tends to focus on earlier periods, when the aggregate effect of El Niño shocks on commodity prices is inflationary.

In Section 7, we put forward a number of possible explanations which may help rationalize these results. For instance, we discuss that the decrease in euro-area consumer prices during El Niño events in the most recent period can be attributed to factors such as the reform of the EU Common Agricultural Policy, which reduced the volatility of EU food commodity prices, and the increased trade diversification and integration, which mitigated the impact of international commodity price shocks on the euro area. Additionally, the substitution of food commodities, where providers switch to less expensive sources not affected by El Niño, could help explain the disinflationary effect of El Niño on euro-area consumer prices in the most recent period. We remind the reader that the primary aim of this paper is to provide empirical evidence on the impact of El Niño and La Niña on euro-area prices, rather than to offer comprehensive explanations for all our findings. Nevertheless, we present a range of possible channels that may contribute to understanding the results.

Our results are related to the recent study by Peersman (2022), who investigates the causal effects of international food commodity price shocks on euro-area inflation by means of instrumental structural vector autoregressions. He uses a series of unanticipated harvest shocks as an instrument to identify exogenous international food

⁷In the 1980s, El Niño shocks increased both international and EU commodity prices, while La Niña shocks increased EU commodity prices but had no significant effects on international commodity prices.

commodity price shocks. This instrument is estimated as the residual from a harvest equation that controls for a number of possible determinants of food production, including global weather phenomena—represented by ENSO events.⁸ He finds that exogenous food commodity price shocks (a 1 percent increase in international commodity prices) have a strong impact on international consumer prices (1.18 percentage points (pp) after eight quarters) and on European Union commodity prices (1.07 pp). This inflationary result is in line with our finding for La Niña, over a similar sample. However, we find a deflationary impact of El Niño on the same sample. In any case, our results are not directly comparable to those of Peersman (2022) for a number of reasons: (i) we consider the direct impact of ENSO shocks, distinguishing between El Niño and La Niña shocks, while Peersman (2022) uses shocks to global harvests of corn, wheat, rice, and soybeans outside of Europe and unrelated to global economic developments to identify exogenous international food commodity price shocks; (ii) he focuses on a narrower set of components for food prices in the EU (cereals, meat, and dairy); (iii) he applies a weighting scheme for EU food commodity prices based on production (Eurostat economic accounts for agriculture), whereas we consider a harmonized index of consumer price (HICP) weights.

We contribute to the existing literature along a number of dimensions. First, we study the impact of El Niño and La Niña on euro-area inflation using both euro-area consumer prices (which are the target of the ECB's monetary policy) and European Union (EU) food commodity prices, given the relevance of EU commodity prices for assessing the pass-through from international food commodity prices to the relevant prices within the euro area (for a detailed discussion, see Ferrucci, Jiménez-Rodríguez, and Onorante 2012 and Borrallo et al. 2026). This is key, since the international and EU food commodity baskets differ both in terms of their composition and weight. This implies that changes in global prices for specific commodities induced by El Niño should not translate one-to-one to the euro-area

⁸In the harvest equation, a global food production index excluding European harvests is regressed on a set of variables reflecting a trend, ENSO shocks, and a vector of control variables that could affect global food production. The variables measuring ENSO shocks represent global weather phenomena that may simultaneously affect European and non-European harvests.

commodity prices, because that commodity may not be as relevant for the European region.

Second, we rely on local projections à la Jordà (2005), in line with the most recent literature on the economic effects of ENSO events (e.g., Cashin, Mohaddes, and Raissi 2017; Atems and Sardar 2021; Adolfsen and Lappe 2023). In this context, most of the existing evidence employs a top-down approach (estimating the impact of El Niño/La Niña on aggregate prices directly) or a disaggregated approach (the impact of El Niño/La Niña on specific components). Our methodological contribution is that we rely on a bottom-up approach, in which disaggregated effects on specific food components are correctly factored in to assess the overall effect of El Niño on aggregated prices. This is important because direct effects on prices of specific components may offset each other, and this may bias top-down results (Ferrucci, Jiménez-Rodríguez, and Onorante 2012).

Third, although our main results refer to the period from 1997 onward, i.e., since the existence of the euro area, we complement them by constructing a new database of EU prices spanning from 1972 onward, which allows us to test the effect of ENSO events on a larger sample that includes quite extreme ENSO episodes dating back to the 1970s.

Fourth, we show that our methodology allows for the replication of most of the existing evidence focusing on international food commodity prices.⁹ This reassures us about the validity of our method and is an indirect test to show the robustness of our main results based on the euro area.

The rest of the paper is organized as follows. Section 2 provides a review of the relevant literature. In Section 3, we describe the data. Section 4 outlines the empirical framework, and in Section 5 we present our results. Finally, in Section 6 we present a number of robustness exercises, while in Section 7 we discuss a number of possible channels that can explain our findings. Section 8 concludes with final observations and remarks.

⁹International food commodity prices are measured by the World Bank's Commodity Price Index, in the World Bank (WB) Pink Sheet Database.

2. Literature Review

The literature on the impact of El Niño events suggests heterogeneous GDP growth and inflation responses to El Niño shocks depending on the country. As regards economic growth, tropical countries experience economic losses (mostly in the agriculture, infrastructure, and energy sectors) to the extent that they are directly affected by weather anomalies. By contrast, El Niño shocks benefit economic activity in countries in temperate regions, which are not directly exposed to these weather-related anomalies. The latter can be explained by, among other factors, trade and aggregate demand resulting from weather disruptions (Brunner 2002; Laosuthi and Selover 2007; Martín 2016; Cashin, Mohaddes, and Raissi 2017; Smith and Ubilava 2017; Beltran, Hadzi-Vaskov, and Stepanov 2023). Whenever asymmetries are considered, GDP growth responds negatively to El Niño shocks, while La Niña shocks do not affect economic activity (Smith and Ubilava 2017). This literature is summarized in Table C.1 in Appendix C. The inflationary effects of El Niño shocks are also heterogeneous across countries (e.g., Brunner 2002; Laosuthi and Selover 2007; Ubilava and Holt 2013; Martín 2016; Cashin, Mohaddes, and Raissi 2017; Abril-Salcedo, Melo-Velandia, and Parra-Amado 2020; Atems and Sardar 2021; Adolfsen and Lappe 2023; Romero and Saldarriaga 2024) and are summarized in Table C.2 in Appendix C. Thus, a disaggregated perspective is crucial to correctly gauge the transmission of weather shocks to aggregate prices, as such heterogeneity can stem from the following several dimensions.

First, from a global perspective, specific food commodity prices respond differently to the same shock: for instance, palm and coconut oil prices increase after El Niño shocks, whereas soybeans and maize are not affected (Atems and Sardar 2021).¹⁰ Second, while weather anomalies are, in principle, a local phenomenon (they directly affect production and the prices of food commodities grown in specific regions/countries), in some cases their effects may spread globally

¹⁰When accounting for asymmetry in the weather anomaly, La Niña shocks increase aggregate food commodity prices while El Niño shocks have ambiguous results, with some food commodities being positively affected and others being negatively affected (Atems and Sardar 2021).


to prices of highly substitutable commodities (as in the case of oils and cereals). Hence, the transmission depends on the degree of substitutability of the affected commodity goods (Ubilava and Holt 2013; Ubilava 2023). Third, responses differ according to the region (tropical regions are more affected, as opposed to temperate regions) and the country. For example, a shock may be transmitted to local food commodity prices, depending on the degree of diversification of the country and its size (Laosuthi and Selover 2007), or the shock may be diluted in temperate countries since food commodities account for only a fraction of the cost of finished goods (Brunner 2002). Fourth, the transmission to consumer prices depends, among other factors, on the share of the exposed food that is present in a country’s consumer price index (CPI) basket (Cashin, Mohaddes, and Raissi 2017). In total, the existing evidence on the transmission to CPI inflation in Europe amounts to 0.07 percentage points (Cashin, Mohaddes, and Raissi 2017) or has been found to have only a marginally significant effect (for the G7 countries) (Brunner 2002).

3. Data

Our variable of interest is the Oceanic El Niño Index (ONI), the main indicator for monitoring the seasonal El Niño–Southern Oscillation climate pattern, from the U.S. National Centers for Environmental Information. It indicates the extent to which the three-month average sea surface temperature of a region of the Pacific Ocean (120°W–170°W) is significantly warmer (El Niño $> 0.5^{\circ}\text{C}$) or colder (La Niña $< -0.5^{\circ}\text{C}$) than a historical average. Five consecutive index values above (below) the threshold of $+0.5^{\circ}\text{C}$ (-0.5°C) are considered to represent an El Niño (La Niña) episode (see Figure 2).

To measure international food commodity, we use the global food commodity price index, as available from the World Bank’s Commodity Price “Pink Sheet.” Database, which is available at a monthly frequency from 1960 onward, as well as the price of its food commodity components: cereals, oils and meals, and other food and beverages. In the aggregate index the latter components are weighted according to the export value of each item relative to the total exports of low-income countries. In this paper all indices are converted from U.S. dollars to euros. Further information on the data composition and sources is available in Appendix B.

Figure 2. Episodes of El Niño and La Niña



Note: The Oceanic Niño Index shows the three-month average sea surface temperature in tropical Pacific Ocean compared to the 1981–2010 average.

As for food commodity prices for the euro area, we consider the food commodity prices of the European Union, as the euro-area countries trade within the single market of the EU for products. This is an important contribution to the literature, which often studies the impact of El Niño events on international food commodity prices, based on the criterion that El Niño is a global phenomenon that affects prices worldwide. By contrast, to the extent that our object of analysis is euro-area prices, we explicitly consider EU food commodities. This allows us to enhance the existing literature in two directions. First, we can take into account the role of the EU’s Common Agricultural Policy (CAP). The CAP smooths the transmission of shocks from international food commodity prices to European farm-gate prices. Due to this policy, European farm-gate prices have been less volatile than international food commodity prices. Second, there is an important difference between aggregate EU and international food commodity prices, which stems from changes in terms of (i) composition—different food items entering into the computation of the aggregate index—as well as in terms of (ii) weights. As shown in Table 1, both the composition of the index and its weights can differ substantially, especially the “vegetable oils” component.

**Table 1. International vs.
EU Food Commodity Price Index**

Group	Item	WB Weight	EU Weight
Food		100.0	100.0
Cereals		28.2	25.0
	Rice	8.5	
	Wheat/Durum	7.1	8.3
	Maize	11.5	
	Barley	1.2	8.3
	Rye		8.3
Vegetable Oils		40.8	5.0
	Soybeans	10.1	
	Soybean Oil	5.3	1.3
	Soybean Meal	10.7	
	Palm Oil	12.3	
	Coconut Oil	1.3	
	Groundnut Oil	1.1	
	Olive Oil		1.3
	Rapeseed Oil		1.3
	Sunflower Oil		1.3
Dairy and Eggs		0.0	21.0
	Raw Milk		5.3
	Skim Milk Powder		5.3
	Edam Cheese		5.3
	Eggs		5.3
Other Food		31.0	49.0
	Sugar	9.8	9.0
	Bananas	4.9	
	Meat, Beef (Cows, Steers)	6.8	18.0
	Meat, Chicken	6.0	9.0
	Meat, Pork		9.0
	Oranges	3.6	
	Coffee		4.0
Note: This table compares the composition and weights of the international food commodity price index, as available from the World Bank Pink Sheet Database, and the European Union food commodity price index (own calculations).			

We construct two data sets to test the impact of both El Niño and La Niña events on euro-area inflation in different time periods. The baseline sample starts from 1997, which coincides with the establishment of the euro area.¹¹ This is our natural first choice, since we are interested in the impact of ENSO events on euro-area inflation. In addition, we construct a larger sample that tracks euro-area prices from 1970, covering a period in which the euro area did not yet exist. This larger sample gives us an additional tool to check the robustness of our results along three directions: first, comparing with the existing literature, which typically considers periods including at least the 1980s; second, including more ENSO events; third, discussing the implications of the CAP on euro-area prices. In the rest of this section, we describe in detail which variables we consider and how we construct the relevant series of interest.

3.1 Baseline Sample: European Food Prices since 1997

This database contains monthly food prices at different levels of the food production process (food consumer, food producer, and food commodity prices) from January 1997 to December 2023 for the euro area and for the four largest European countries (France, Germany, Italy, Spain). Food consumer prices and food producer prices are obtained from Eurostat, while for food commodity prices we rely on the European Commission's Directorate-General for Agriculture and Rural Development (DG Agri) database.

Specifically, we compile data on food prices for six food groups: cereals, dairy, fats and oils, meat, sugar, and coffee, representing close to 70 percent of the food consumption basket. We exclude fruit and vegetable prices as they are highly volatile, as well as fish prices as there is no reference for its commodity price at the European Union level.¹² In any case, despite data unavailability, the

¹¹The euro zone was born with its first 11 member states on January 1, 1999. However, euro-area data are gathered since 1997.

¹²Commodity prices are unavailable in the European Union for fish, and there are very few items for fruits and vegetables, which raises concerns about the representativeness of the resulting food category. This is the key reason why we considered their exclusion from the analysis, as we cannot guarantee the same quality standards in the data, and in some cases we cannot even run the same model.

exclusion of fish, vegetables, and fruits should not affect our results significantly, as the EU imports fish from diversified sources and is largely self-sufficient in vegetables and nearly independent in fruits. In particular, the EU is a net importer of fish, as it produces only 60 percent of its consumption needs domestically, but the import part is sourced from highly diversified providers.¹³ The EU is self-sufficient to produce vegetables, and close to independent for fruits, as just 14 percent of its consumption depends on imports (Cuadro-Sáez, Pérez, and Viani 2025).

Food consumer prices and food producer prices are obtained from Eurostat, using the corresponding category for each group (see Table A.1 in Appendix A). We then consider the European Union’s corresponding food commodity prices available from the European Commission’s DG Agri. This database provides highly detailed indices, allowing us to construct an index for each food group commodity as the simple average of its components. For instance, the cereal food commodity price index is the mean price of feed barley, malting barley, feed maize, bread wheat, feed wheat, durum wheat, feed rye, and bread rye (refer to Table A.1 for the composition and correspondences among food price indices). When data are not available for commodity prices, we use international commodity prices instead.¹⁴

As for control variables, we include as a proxy for economic activity the growth in the industrial production index (IPI) for the euro area, as available from Eurostat. We also control for oil price and fertilizer price indices as provided by the World Bank’s Pink Sheet, which we convert into euros using the spot exchange rate. Both variables are defined at the global level.

3.2 Larger Sample: European Food Prices since 1972

We construct a novel database from January 1971 onward. Data on consumer prices are not available for any of the six groups, so we

¹³Despite the significant weight of Norway, which contributes 35 percent of imports.

¹⁴ This relates to only a few components: prices of rapeseed oil, sunflower oil, and soybean oils (which belong to the fats and oils category) are taken from the International Monetary Fund (IMF) database; the sugar price is provided by the WB Pink Sheet; the coffee price is from the International Coffee Organization (see Table A.1).

focus on the available aggregates: “headline” and “food and non-alcoholic beverages” consumer prices. In contrast, food commodity prices for the European region are available for this longer period of time. To extend these series backwards, we combine monthly information available from Eurostat (for the 1971–95 period) and DG Agri from 1996 onward.¹⁵ Specifically, for data before 1996 we take the components that appear both in the DG Agri database and Eurostat and whose monthly growths have a correlation of 65 percent or more in the 1991–94 period, and extend DG Agri price levels backwards using the monthly growth implied in Eurostat data. Then, after 1996 we maintain the indices used in the previous section (see Table A.2 in Appendix A).

4. Methodology

We estimate the impact of an intensification of El Niño on prices by using the local projections (LPs) method of Jordà (2005), in line with the most recent literature (e.g., Atems and Sardar 2021 and Adolfsen and Lappe 2023). The LPs estimator is attractive since it directly projects an outcome at the future horizon h on current covariates and does not require specification and estimation of the unknown true underlying multivariate system. For this reason, we believe that in our setting LPs is the most suitable methodology.

We consider the following specification:

$$\begin{aligned} \pi_{t+h} = & \alpha + \sum_{i=1}^j \beta^h \cdot \pi_{t-i} + \gamma^h \cdot \text{niño}_t + \sum_{i=0}^j \nu^h \cdot \text{fuel}_{t-i} \\ & + \sum_{i=0}^j \delta^h \cdot \text{fert}_{t-i} + \sum_{i=0}^j \kappa^h \cdot \text{IPI}_{t-i} + \epsilon_t, \end{aligned} \quad (1)$$

where $h = 0, \dots, 18$, $j = 6$, and h and i refer to months. The variable of interest, El Niño, equals the Oceanic El Niño Index (ONI) during an El Niño event and takes the value 0 otherwise (that is,

¹⁵We use monthly data from Eurostat from 1971 to December 1995, and we link these with monthly data from DG Agri from January 1996 onward for all available components since 1971, except for egg prices, which are available from 1997 onward.

during La Niña events and when neither of the two take place). When we consider La Niña, we construct the same variable for these events, taking the absolute value of the resulting indicator.¹⁶ Control variables include oil prices ($fuel_t$), fertilizer prices ($fert_t$), and the industrial production index (IPI_t) as a proxy for economic activity. All variables in the model refer to log month-on-month changes.¹⁷ The dependent variable is a measure of the euro-area price inflation, which can be defined as consumer prices or commodity prices, and can be set at the aggregate level or disaggregated at a more granular level, depending on the specification. Additionally, we suppose that industrial production and fuel and fertilizer prices are exogenous to food commodity prices.

The shock considered in the analysis is the effect of an intensification of an El Niño event, i.e., an increase of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño event is already taking place. Since the *niño* variable takes the ONI value when an El Niño event is taking place and 0 otherwise, the intensification of an El Niño event implies a rise of γ percentage points in price growth. This interpretation applies also for La Niña shocks, since intensifications of this event enter with the absolute value in the equation. Note, the *niño* variable enters into the equation contemporaneously, as the estimation structure of local projections allows us to calculate how it affects future values of the target variable through the impulse response function (IRF).

For each specification, we choose the optimal lag structure according to the following two-step strategy.¹⁸ First, we estimate the model including the same number of lags—from one to six lags—for all control variables except for the El Niño variable, which enters only contemporaneously. For each specification, we compute the Hannan-Quinn information criterion, and then we choose the number of lags that corresponds to the model that performs best according to the Hannan-Quinn criterion. In a second step, for each

¹⁶In Appendices J and K we show that our results are robust to considering them simultaneously.

¹⁷Oil, fertilizer, euro-area industrial production, and food commodity and consumer prices are seasonally adjusted by using X-13 ARIMA-SEATS.

¹⁸For all cases, in order to find the optimum specification we apply this procedure to the first-horizon regression of the LP exercise and then keep the resulting specification for all horizons.

model we exclude the non-significant lags from the regression (aside from the contemporaneous and the first lag, which always remain).¹⁹ Finally, our linear projection estimations account for heteroskedasticity by using Huber-White corrected standard errors, as suggested in Montiel Olea and Plagborg-Møller (2021).

We run our estimations from different perspectives. First, we run the analysis at the disaggregated level and provide the corresponding IRF of each of the six specific food groups. Second, we estimate the model on the aggregate price indices, which are weighted averages of the group-specific food components. This top-down approach is what is commonly reported in the literature. This method may lead to composition bias since direct effects on prices of specific components may be offsetting each other, as pointed out by Ferrucci, Jiménez-Rodríguez, and Onorante (2012). We solve this issue by means of a bottom-up approach. That is, we first estimate the response of each specific component to an El Niño shock and then group responses to an aggregate level (Ferrucci, Jiménez-Rodríguez, and Onorante 2012). In the aggregation, the median responses are aggregated by means of a weighted average. The construction of the confidence bands of the aggregate IRFs in the bottom-up approach is based on a block bootstrap procedure with fixed regressors. We jointly resample the estimated residuals of the 12 models (one model for each of the six food commodity price components plus one model for each of the food consumer price components) via a block bootstrap and combine the resampled residuals with the fixed regressors. This gives us a bootstrapped series of commodities and consumer food price values on which we reestimate the models and the IRFs. We repeat the bootstrap procedure 1,000 times and construct the

¹⁹For all cases, the optimum number of lags is one lag, except for the models of dairy consumer prices (two lags), coffee consumer prices (two lags), and sugar consumer prices (three lags). Hence, step two of the procedure applies only to these latter specifications. It turns out that for the model of dairy consumer prices, we keep the second lag of the dependent variable and the economic activity variable and drop the second lag of oil and fertilizer prices. For the model of coffee consumer prices, we keep only the second lag of the dependent variable and drop the second lag for all other controls; for the model of sugar consumer prices, we include up to the third lag of the dependent variable and the third lag of the activity indicator, while the rest of the controls enter with the first lag only.

respective confidence bands for the aggregated IRFs by aggregating the IRFs for each individual bootstrap run and then extracting the corresponding percentiles. This allows us to compute confidence bands that take into account the cross-sectional correlation of the error terms across distinct food commodity and consumer price components. For a general description of bootstrapping methods, see Kreiss and Lahiri (2012).

5. Results

All figures and the discussion below should be read as referring to the effects of shocks defined as an increase of 1°C in the ocean surface temperature during an El Niño event (ONI) and a decrease of 1°C in the ocean surface temperature during a La Niña event (ONI)—thus, an intensification of either an El Niño or La Niña episode.


5.1 *Impact on International Food Commodity Prices*

Results of both El Niño and La Niña shocks on international food commodity prices over the longer sample period (1972–2023) are in line with the existing evidence (e.g., Brunner 2002; Cashin, Mohaddes, and Raissi 2017; Adolfsen and Lappe 2023) that currently shapes discussions about the risks implied by these climate-related phenomena.

Our results show that the intensification of an El Niño event leads to a significant increase in global food commodity price growth of about 9 percentage points one year after the shock and that this cumulative effect remains significant for at least 18 months after the shock. In contrast, the impact of La Niña is barely significant in the fourth month after the shock, causing an increase in international food commodity prices somewhat below 2 percentage points (see Figure 3) and is otherwise not significant.

Appendix D reports food-group-specific effects of ENSO shocks. For the case of El Niño, most of the spike in the aggregate food commodity index is due to the response of vegetable oils, and to a lesser extent sugar and cereals. Results for El Niño are broadly consistent with the previous literature. For example, Smith and Ubilava (2017) document a decrease in global production, which then implies an increase in the price. Other studies find a positive impact on the

Figure 3. International Commodity Price Responses to an Intensification of El Niño/La Niña



Note: Monthly cumulative response of international commodity prices to an El Niño/La Niña shock along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. Results are based on linear projection estimations. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1 °C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.

price of cereals (Brunner 2002; Laosuthi and Selover 2007; Ubilava 2017; Adolfsen and Lappe 2023), on the global prices of vegetable oils (Ubilava and Holt 2013), and on the price of sugar (Cashin, Mohaddes, and Raissi 2017). However, for a few specific items we find it hard to reconcile our results with previous ones, essentially because the existing evidence is not clear-cut, i.e., Brunner (2002), Martín (2016), and Cashin, Mohaddes, and Raissi (2017) show a positive impact of El Niño on coffee, while Ubilava (2011) demonstrates that the impact of El Niño depends on the type of coffee, finding a positive impact of El Niño on robusta coffee beans and a negative impact on arabica coffee beans.


Given our aim of studying the impact of the intensification of ENSO events on euro-area inflation, we split the sample into two periods. The first spans from 1972 to the Common Agricultural Policy reform in 1992, known as MacSharry’s reform, by which the EU moves from indirect support to farmers via guaranteed prices to direct support via income transfers.²⁰ The second period starts in 1997, shortly before the launch of the euro.

Splitting the sample reveals two striking results (see Figure 4). First, the impact of El Niño on international commodity prices is very intense in the former period (1972–92), i.e., an intensification of El Niño causes an increase of about 20 percentage points in international food commodity prices after 12 months. Second, and more importantly, this impact completely fades away when considering the more recent euro-area sample (1997–2023). One possible explanation is that during the former period, which includes the 1980s, four strong El Niño events occurred consecutively in a closely spaced manner. In contrast, the most recent period features strong El Niño events that occurred with more time between each occurrence. Section 7 presents a discussion of further additional economic explanations.

In the case of La Niña, the sample split shows that the former period is consistent with existing findings along two dimensions:

²⁰In the guaranteed prices system (before 1992), excessively high food commodity prices generate an excess of production (sustained also by excessive stocks and export subsidies) which, in turn, creates negative consequences such as the fall of farm income and soaring budget expenditure, while consumption increases at a slower pace than production (Garzon 2006).

Figure 4. International Commodity Price Responses to an Intensification of El Niño/La Niña: Split Sample



Note: Monthly cumulative response of international commodity prices to an El Niño/La Niña shock. Full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations using a bottom-up approach. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.


(i) La Niña tends to exert a softer effect on prices than El Niño, and (ii) if anything, La Niña slightly reduces international commodity prices. In any case, the maximum impact takes place around the fifth month after the shock, with an increase in international food commodity prices of about 4 percentage points, which is still well below the estimated impacts of El Niño on food commodity prices. This confirms that La Niña has softer effects than El Niño. The positive effect found in the most recent period might be explained by the fact that the 2000–12 period featured four strong La Niña events that happened consecutively in a closely spaced manner.

5.2 Impact on EU Food Commodity Prices

Turning to the impact on EU food commodity prices, and maintaining the sample divided into two periods, results show that in 1972–92 El Niño significantly affected EU food commodity prices by about 7 percentage points 12 months after the shock. This effect vanishes when considering the most recent period (see Figure 5). Notably, in the most recent period, the effect of El Niño on food prices turns negative, although the coefficients are small and only marginally significant. Disaggregated responses by food groups reveal that El Niño has a clear disinflationary effect on some commodities, such as meat and cereals, which is offset by the responses of other food groups. Moreover, the disinflationary impact of El Niño is confirmed and becomes even stronger when taking into account the pass-through along the value chain (see Appendix F). While we do not aim to explain the reaction of each food group, in the rest of this section and in Section 7 we propose some possible explanations that may help rationalize the disinflationary effect of El Niño on EU food prices. A more detailed discussion on these disaggregated effects is provided in Appendix E.

Note that during the first period, food commodity prices were guaranteed in the European Union, such that agricultural producers had an incentive to overproduce. As long as El Niño adversely affected international production, European producers (supported by export subsidies) had an incentive to export their surpluses at international prices. The CAP reform in 1992 entailed direct support to farmers via income. The removal of the incentives for EU farmers to overproduce and export as of 1992 is consistent with the

Figure 5. EU Commodity Price Responses to an Intensification of El Niño/La Niña: Split Sample



Note: Monthly cumulative response of EU food commodity prices to an El Niño/La Niña shock. Full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations using a bottom-up approach. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.

hypothesis that El Niño no longer had any effect on EU commodity prices.


By contrast, La Niña induces a moderated and short-lived (up to five months after the shock) inflationary effect on EU food commodity prices, an effect that is robust over time. This apparent inconsistency between the response of international (not significant in the short term and turning negative afterward) and EU commodity prices (positive) in the case of a La Niña shock in the former period (until 1992) can be explained looking at the disaggregated effects (see Figure D.2 in Appendix D and Figure H.2 in Appendix H). In both price indicators, the positive impact comes from meat and sugar, while the negative one is due to vegetable oils. What drives the difference is the different weighting schemes used to aggregate food components: vegetable oils have a large weight in the international commodity index and a small weight the EU one. Conversely, meat counts much more in the EU commodity weighting scheme than in the international scheme (see Table 1).

5.3 *Consumer Price Results*

In this section, we focus on the target of our analysis, i.e., the impact of ENSO events on food inflation in the euro area, for which we consider the euro-area sample since 1997. We consider a restricted definition of the HICP food price index consistent with our EU commodity price index, in which we retain six food groups that amount to 70 percent of the food consumer basket (see Figure A.1 in Appendix A).²¹ In particular, we consider cereals, dairy, fats and oils, meat, sugar, and coffee and leave aside the other food groups (mineral water and juices, fruits and vegetables, and fish) due to their high volatility and the lack of consistent price data available. This approach results in a negative and significant impact of an intensification of El Niño of about 0.6–0.7 (top-down/bottom-up) percentage points on food inflation 12 months after the shock, as the impact on commodity prices passes through to food inflation and to headline inflation. The impact is positive and more pronounced in the case of

²¹Results based on the aggregate HICP food price index are in line with our findings for our restrictive definition of the HICP food price index and are available to the reader upon request.

Figure 6. Response of Euro-Area Food Consumer Price Inflation to an Intensification of El Niño/La Niña



Note: Monthly cumulative response of the euro-area food HICP based on our own six food group aggregate to an El Niño shock (left panels) and La Niña shock (right panels) along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. In the third row, full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.

the intensification of a La Niña event (1–1.1 percentage points) (see Figure 6).

The impact on food inflation passes through to headline inflation in a significant way. Contrary to common belief, we find that an intensification of El Niño weakly reduces euro-area inflation by

0.3 percentage points. Besides, the effect of an intensification of La Niña events is positive and of about 0.6 percentage points in cumulative terms 12 months after the shock (see Figure 7). These two results are crucial in the current circumstances as (i) El Niño might not have been pushing inflation upward as believed and, more importantly, (ii) one should take into account the risks to inflation due to an intensification of La Niña events.²²

6. Robustness


We provide a number of robustness exercises. First, we check the robustness of the baseline results by expanding our model to follow the value chain perspective of Ferrucci, Jiménez-Rodríguez, and Onorante (2012). We do so by adding at least one lag of food producer prices and of food consumer prices in the food commodity equation, and at least the first lag of food producer prices and food commodity prices in the model for food consumer prices (Equation (1)). Results are reported in Appendix F and are very similar to our main results.

Second, we carry out the analysis at the country level for the most important euro-area countries: Italy, France, Germany, Spain (see Appendix G). The responses of food inflation to intensifications of El Niño and La Niña for each of those countries are quite similar to the responses of the euro area as a whole.

Third, we reestimate the analysis for consumer prices using the larger sample (from the 1970s), for the euro area (see Appendix H). Unfortunately, data for each of the six consumer prices were not available from 1971, so we estimate the effects of El Niño and La Niña shocks on the food component of the HICP, the headline component, and commodity prices. The response of food inflation to both shocks was similar to the top-down and bottom-up results of the main exercise. However, headline inflation in the euro area does

²²On a conjunctural note, at the time of writing, the El Niño–Southern Oscillation is likely to turn into a La Niña phase this summer: there is a 65 percent chance that La Niña is declared between July and September 2024 and persists into the Northern Hemisphere winter of 2024–25 (85 percent chance during November–January), according to the National Oceanic and Atmospheric Administration’s (NOAA’s) Climate Prediction Center. See the “ENSO: Recent Evolution, Current Status and Predictions” report of June 24, 2024.

Figure 7. Response of Euro-Area Headline Consumer Price Inflation to an Intensification of El Niño/La Niña



Note: Monthly cumulative response of the euro-area headline HICP to an El Niño/La Niña shock along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. Results are based on linear projection estimations. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.

not react significantly to El Niño, although it still reacts significantly to La Niña, and commodity prices increase significantly after both phenomena.

Fourth, as an additional robustness analysis, we run the baseline model for the international food commodity prices adding, on one hand, the one-month change in the dollar-to-euro exchange rate and, on the other hand, changing the number of lags for food commodity, fertilizer, and oil prices and activity to 12. The resulting figures are shown in Appendix I. For El Niño shocks, adding the exchange rate to the baseline model does not change the results, but including more lags reduces their effects on food commodity price growth by around 20 percent relative to the baseline model. However, for La Niña, adding the exchange rate makes the short-run positive effect of the phenomenon on commodity price growth not statistically significant, while no meaningful change happens when adding 12 lags to the baseline model.


Fifth, we included both El Niño and La Niña variables in the same regression model, rather than analyzing them separately. This approach acknowledges that El Niño and La Niña are interrelated phases of the ENSO climate phenomenon. Our findings indicate that the baseline results are robust when both events are treated simultaneously (see Appendix J).

Finally, we include in the same regression El Niño and La Niña, exploiting all variation in the ONI variable. Hence, we consider the complete variation of the ONI variable, not only the excess from the normal value. This allows us to take into account other dimensions of the ONI events, such as duration and intensity. Results are reported in Appendix K and are in line with our baseline conclusions.

7. Discussion

With respect to food commodity prices, we show that there are nuances that have to do with composition effects (type of commodity) and sample periods (with more recent decades presenting smaller effects than the former period observed). With respect to consumer price inflation in the euro area, we show that it is crucial to account for the role of domestic agricultural policies (for the euro area, the CAP) and trade patterns. In particular, focusing on the euro-area

Figure 8. Channels Layout



Note: For mitigation policies, see Section 7.2; for the concentration of extreme events, see Section 7.1; for the CAP, see Section 7.4; for trade diversification, see Section 7.5; for more integrated commodity markets, see Section 7.3.

(EA) sample, an intensification of El Niño decreases EU food commodity prices by 1 pp, which in turn decreases EA headline inflation by about 0.3 pp after 12 months.

As for La Niña shocks, these increase EU food commodity prices by 1.8 pp, which passes through to EA headline inflation (+0.6 pp after 12 months). While the impact of the intensification of La Niña on EU food commodity prices is stable over time,²³ the impact of El Niño shocks on international commodity prices changes over time. These are positive and large when including the 1970s and 1980s in the analysis, whereas they fade away (for international commodity prices) or even switch signs (for EU commodity prices) when restricting to the EA sample.

Our findings can be explained by the combination of a number of factors. For the convenience of the reader, we represent these arguments in the following diagram (see Figure 8) and explain them one by one.

²³The effect of an intensification of La Niña on international food prices is negative in the '70s–'80s and becomes inflationary in the most recent period. We do not have an explanation for that.

7.1 Concentration of Extreme Phenomena in the 1980s

A possible explanation for the smaller impact of El Niño shocks in the most recent sample might be a change in the distribution of ENSO events. In Table 2, we report descriptive statistics for all ENSO events included in this analysis. Strong ENSO events are shown in bold. For El Niño, we note that in the former period, and especially during the 1980s, there was an extreme concentration of strong El Niño phenomena (three strong El Niño events in a row within the decade from 1982 to 1992). Moreover, these events are associated with a striking variation in commodity prices. These episodes took place in a time when food inflation volatility was high and market integration was not as developed as it later became. As a consequence, the combination of a high concentration of strong El Niño events in the 1980s and the strong food price variation in that decade may have driven the results for the former sample, and hence in part explain our findings.

7.2 Mitigation Policies and the Financialization of Food Commodities

Early warning systems allow farmers to frontload or delay the planting and harvesting of crops when the formation of an El Niño event is detected, thereby mitigating potential losses (see, for instance, Adams et al. 2003; Paz et al. 2007; Raj, Kumar, and Ramesh 2020). This channel could help explain our results. Indeed, fully fledged policies for anticipating, responding to, and mitigating the effects of the ENSO phases have been called for by the Food and Agriculture Organization (FAO). For instance, FAO (2023) points to policies that help fishers protect their boats ahead of storms or reinforced river embankments ahead of floods. The organization also calls for responding to harvest damage by distributing drought-tolerant seeds to rainfed farmers and protecting livestock health, among other strategies. The FAO also emphasizes capitalizing on the positive spillover effects of El Niño and offsetting its losses, for instance, by supplying seeds to flood-affected farmers so they can plant and regain a harvest as flood waters recede. A third step the FAO calls for is the fast delivery of the most-needed products for farmers in case of devastation, such as veterinary medicines, seeds,

Table 2. Characteristics of ENSO Phenomena

El Niño				
Start	End	Length [†]	EU Comm. [§]	Int. Comm. [¶]
May-72	Mar-73	12	13.9	37.8
Sep-76	Jan-77	6	4.9	3.2
Sep-77	Dec-77	5	2.7	9.4
Oct-79	Jan-80	5	3.8	1.1
Apr-82	Jun-83	16	9.4	11.9
Sep-86	Feb-88	19	-1.5	12.7
May-91	Jun-92	15	2.0	-4.8
Sep-94	Mar-95	8	-0.7	-2.0
May-97	May-98	14	-9.9	-2.8
Jun-02	Feb-03	10	5.0	-4.2
Jul-04	Feb-05	9	0.6	-6.4
Sep-06	Jun-07	6	1.7	8.4
Jul-09	Mar-10	10	4.6	4.2
Oct-14	Apr-16	20	-5.6	-0.7
Sep-18	Jun-19	11	1.2	6.2
May-23	Dec-23	9	-5.4	-5.9
Avg. All Events		10.9	1.7	4.3
Avg. Strong Events (*)		14.4	1.0	6.6
La Niña				
Start	End	Length [†]	EU Comm. [§]	Int. Comm. [¶]
Jan-71	Jan-72	14	11.1	-6.1
May-73	Jul-74	16	36.6	38.5
Oct-74	Apr-76	20	-16.2	-33.8
Sep-83	Jan-84	6	-0.6	1.8
Oct-84	Aug-85	12	-7.1	-20.6
May-88	May-89	14	16.8	28.1
Aug-95	Mar-96	9	-4.5	8.8
Jul-98	Feb-01	33	3.5	0.1
Nov-05	Mar-06	6	9.7	4.6
Jun-07	Jun-08	14	13.6	38.7
Nov-08	Mar-09	6	-4.8	-0.1

(continued)

Table 2. (Continued)

La Niña				
Start	End	Length [†]	EU Comm. [§]	Int. Comm. [¶]
Jun-10	May-11	13	23.6	17.9
Jul-11	Apr-12	11	−4.7	7.8
Aug-16	Dec-16	6	5.2	4.9
Oct-17	Apr-18	8	−7.2	3.9
Aug-20	May-21	11	19.0	36.3
Aug-21	Jan-23	19	28.8	15.6
Avg. All Events		12.8	7.2	8.6
Avg. Strong Events (*)		18.3	14.9	17.9
Note: The table shows the month and year in which events of El Niño (La Niña) start and end, the number of months for which these events took place, and the change in both the European and international food commodity prices during each event. Rows in bold refer to strong El Niño and La Niña events. A strong El Niño (La Niña) is defined as an event where the average sea surface temperature in the equatorial Pacific is at least 1.5°C warmer (colder) than normal. [†]Length: Number of months since the start of the event. [§]EU commodity price. We report the rate of change between the final month of the event and the initial month of the event. [¶]International commodity price. We report the rate of change between the final month of the event and the initial month of the event. (*) We compute the average considering only strong events, which are listed in bold in the table.				

and water bladders, while providing cash to severely affected families to meet their most immediate needs.

Another factor that might have mitigated the impact of the El Niño–Southern Oscillation over the most recent decades is the hedging strategies on food commodity derivative markets that have developed markedly since early 2000 and which have proven to be effective in lowering risks (see, for instance, Jiang and Fortenbery 2019; Ferreira, Tonin, and Alves 2022; Pantoja-Robayo and Rodriguez-Guevara 2023). This is probably related to the more intense participation of institutional investors providing further liquidity to the food commodity derivative markets since the late 1990s, known as the financialization of commodity markets, as these traders tend to adopt higher risk-return positions that generate

broader hedging opportunities for commercial investors (ECB 2011; Tang and Xiong 2012; Cheng and Xiong 2014).

We remind the reader that the most likely effect of mitigation policies is to reduce the ONI-induced price increases. However, these policies could also lead to a decrease in prices in certain cases, although this is something we cannot test for. For instance, mitigation policies may enable farmers to anticipate or delay certain cultivation phases in response to an impending El Niño event. If these adjustments result in an increased supply, the resulting harvest could potentially lead to a decrease in food commodity prices.

7.3 The Volatility of Food Commodity Prices in the 1980s


Increased globalization and economic integration may also partly explain why the impact of ENSO phenomena has decreased over time. As is shown in Figure 9, food commodity price volatility has decreased since the middle of the 1980s. Jacks, O'Rourke, and Williamson (2009) shows that world market integration is associated with less commodity price volatility.²⁴ This means that a higher degree of globalization could reduce food price volatility. This argument suggests that food commodity prices have become less reactive to shocks, and hence the effects of ENSO climate phenomena may have become weaker.

7.4 EU's Common Agricultural Policy

The reform of the Common Agricultural Policy (CAP) in 1992, known as MacSharry's reform, may also help explain our findings, i.e., the differences in the inflationary effects of El Niño in different samples. Until 1992, the CAP guaranteed minimum food commodity prices for farm production, and farmers were offered subsidies to export. This combination encouraged excess production and trade out of Europe. Hence, in this setting, when international food commodity prices increased due to El Niño shocks, European farmers had an incentive to export to international markets. As a consequence, food commodity supply decreased in the European market

²⁴The paper studies 300 years of data for commodity prices and links it to periods of high and low globalization.

Figure 9. 24-Month Volatility of International Food Commodity Prices



Note: Own calculations. The figure shows the 24-month standard deviation of the international food commodity index. The horizontal lines show the average from 1973 to 1992 and from 1997 to 2023.

and EU food commodity prices increased due to competition effects. After 1992 this policy changed, such that support to farmers became direct—through rents—and minimum prices and export subsidies were no longer guaranteed.

7.5 Trade Diversification in the Euro Area

The working hypothesis here is that trade patterns may have helped mitigate the effect of ENSO events on EU food commodity prices over time. To shed light on this potential channel, we provide some facts on the evolution of EU trade by food group, in line with our main analysis.²⁵ For this descriptive exercise, we focus on three food groups, namely cereals, meat, and dairy, which together represent 50 percent of the EU consumer food basket.

For each food group and year, we compute the import share, defined as the sum of shares of EU imports coming from countries that are affected or not by ENSO phenomena (either El Niño or La Niña) out of the total EU imports, for the corresponding

²⁵For details on the data and methodology, see Appendix L.

food category. We also compute the Herfindahl-Hirschman Index (HHI) in order to measure the degree of import concentration across trade partners, distinguishing between the import concentration from countries that are affected by ENSO phenomena and the import concentration from other countries that are not directly affected by ENSO phenomena.


The following dynamics stand out. For each food group, the share of imports from ENSO-affected countries has decreased since 2000 to the benefit of the import share from ENSO-neutral countries, which has increased (see Figure 10). This means that since the 2000s, the EU has traded more intensively with ENSO-neutral countries, shifting imports away from ENSO-affected countries. In terms of concentration, for dairy and cereals we observe an increase in the overall concentration index since the 2000s, resulting from both the increased concentration within the market segment of imports from ENSO-neutral countries and the broader gap in import shares to the benefit of this group of countries (see Figure 11).²⁶

Instead, for meat the overall concentration index has remained broadly stable since the 2000s, as a result of two offsetting trends: the increase in concentration within the segment of ENSO-neutral countries and the narrowing of the import share gap to the benefit of this group of countries, which lowers the concentration.

All in all, this descriptive evidence is consistent with the working hypothesis that EU countries may have implemented trade strategies that protect the EU market from the food price volatility implied by ENSO shocks along two dimensions: (i) by shifting the import share from countries affected by ENSO shocks toward countries not affected by these phenomena and (ii) by concentrating within the segment of the import market of countries not affected by ENSO shocks, i.e., importing more from fewer ENSO-neutral countries. This may help explain our findings. However, this conclusion has to be taken with a grain of salt, as trade data cover only the most


²⁶This follows from the fact that the increase in the overall concentration index is coupled, firstly, with the increase in import shares from ENSO-neutral countries and, secondly, with the rise in concentration in this segment of countries in the period when concentration in the ENSO-affected group remains virtually stable.

Figure 10. Import Share by Food Group



Source: BACI International Trade Database at the Product-Level by CEPII (BACI-CEPII).
Note: These figures display import share indices as described in Appendix L. The solid thin blue line depicts the linear trend of the import share from ENSO-affected countries since 2001, to facilitate the reading of the graph.

Figure 11. Import Concentration by Food Group



Source: BACI-CEPII.

Note: These figures display import concentration indices as described in Appendix L.

recent period and, hence, we are extrapolating the trade dynamics backwards.

8. Conclusions

This paper investigates the effect of El Niño and La Niña events on euro-area inflation. We focus on a comprehensive granular database covering both euro-area consumer prices and European Union (EU) food commodity prices. We also construct a new database of EU prices from 1972 onward, allowing us to test the effect of ENSO events on a larger sample. We assess the impact of these climatic events by means of local projections techniques, following the methodology of Jordà (2005).

Our findings based on data from 1997 onward show that an intensification of an El Niño event actually results in a decrease in euro-area headline inflation of approximately 0.3 percentage points 12 months post-shock. In stark contrast, a La Niña event leads to an increase in headline inflation of 0.6 percentage points 12 months after the shock. When we extend our analysis to include data dating back to the 1970s, we identify the 1980s as a pivotal decade demonstrating a significant positive impact of El Niño on both global and European Union food commodity prices. This finding aligns with the existing literature. Interestingly, this trend reversed post-1997.

All in all, our results suggest that El Niño may not have been a key driver of upward inflation pressures in the euro area during the most recent episode of high inflation. Nevertheless, policymakers should remain alert to the inflation risk associated with an intensification of the impending La Niña event.

This paper contributes to the existing literature by filling a gap in the study of the impact of ENSO events on euro-area prices. Our bottom-up approach, which factors in disaggregated effects on specific food components to assess the overall effect of El Niño and La Niña on aggregated prices, provides a more nuanced understanding of these phenomena. Furthermore, our methodology allows us to replicate most of the existing evidence on international food commodity prices, thereby validating our approach and indirectly testing the robustness of our main results based on the euro area.

Appendix A. Food Components across Sources


Table A.1. Food Prices: Definitions and Subcomponents across Food Groups and Production Processes

	Commodity Price	Producer Price	Consumer Price
Cereals	Feed barley, malting barley, feed maize, bread wheat, feed wheat, durum wheat, feed rye, and bread rye	Cereals: manufacture of grain mill products, starches, and starch products	Cereals: bread and cereals
Dairy	Skim milk powder, raw milk, butter, cheddar, edam, and eggs	Manufacture of dairy products	Dairy: milk, cheese, and eggs
Fats and Oils	Rapeseed oil, sunflower oil, soybean oil—the three of them from the IMF—and olive oil	Manufacture of oils and fats	Oils and fats
Meat	Cows, steers, pigs class E, and chicken	Production of meat and poultry meat products	
Sugar	White sugar (from WB Pink Sheet since May 2023)	Manufacture of sugar	Sugar: sugar, jam, honey, syrups, chocolate, and confections
Coffee	Coffee (from the International Coffee Organization)	Processing of tea and coffee	Coffee, tea, and cocoa

Table A.2. Food Commodity Price Indices Before and After 1996

	Components of Commodity Prices Before 1996	Components of Commodity Prices After 1996
Cereals	Feed rye, durum wheat, feed barley, and feed maize (before 1991, feed maize comes from the Pink Sheet)	Feed barley, malting barley, feed maize, bread wheat, feed wheat, durum wheat, feed rye, and bread rye
Dairy	Skim milk powder, raw milk, edam, and eggs	Skim milk powder, raw milk, butter, cheddar, edam, and eggs
Fats and Oils	Soybean oil until 1980, then also rapeseed oil and sunflower oil	Rapeseed oil, sunflower oil, soybean oil—all three from the IMF—and olive oil
Meat	Cows, pigs class E, and chicken	Cows, steers, pigs class E, and chicken
Sugar	Sugar (from the International Financial Statistics up to 1979 and IMF commodity prices from 1980 onward)	White sugar (from WB Pink Sheet since May 2023)
Coffee	Coffee (robusta, from the Pink Sheet)	Coffee (from the International Coffee Organization)

Figure A.1. Correspondence of EU Commodity Prices and Euro-Area Consumer Prices



Note: Own elaboration. N.E.C. means not elsewhere classified.

Appendix B. Detailed Commodity Food Price Data Since 1972

This appendix describes in detail the sources and procedures used to construct the variables since 1972.

Global Food Commodity Prices. Monthly data on global food commodity price indices are obtained from the World Bank’s Pink Sheet. We include in the database the food aggregate as well as each of the food groups and items. Food items are aggregated into groups and the total food aggregate according to the average export value weights in emerging economies in the 2002–04 period. All indices are converted from U.S. dollars to euros.

European Union Food Commodity Prices. To construct this aggregate price index, we proceed as follows:

- We consider monthly data for cereals, meat, and dairy from 1972 to 2023. We concentrate on these three food groups due to data availability restrictions for their components: cereals (rye, durum, feed barley), meat (heifers, chicken, cows, pigs), and dairy and eggs (raw milk, skimmed milk powder, edam cheese, and eggs). We use DG Agri monthly data on food commodity prices from 1991 onward, and we complete the

database with the equivalent Eurostat monthly prices, available on a monthly basis from 1971 to 1995/99 depending on the specific food item. All prices are expressed in euros.


- We include in our database the food items showing a correlation between growth rates from the two data sources higher than 0.65 during the overlapping period (1991–94). For instance, edam cheese prices show the lowest correlation of 0.66 between the growth rates available in Eurostat and DG Agri. Cows’ meat shows the highest (0.95).
- For each food item, we use the Eurostat price index from January 1971 to December 1995 normalized to 100 in January 1971.²⁷ Then, we continue the series, applying to the last observation available from the Eurostat data the corresponding growth rate in DG Agri.
- Each food group price index is calculated as the simple average across its food items’ prices. The aggregate food commodity price for the European Union is the weighted average of the food group price indices, according to the weight of each category in the food consumption basket in 2021 (0.31 for cereals, 0.43 for meat, and 0.26 for dairy and egg products).
- To ensure that this procedure is correct, we compare the resulting food commodity aggregate price index growth rate to the equivalent price series in Eurostat (available at a yearly frequency from 2000 onward). Figure B.1 shows that the two series are very similar over the last two decades, ensuring that our procedure is valid.

Euro-Area Consumer Prices. The Harmonised Index of Consumer Prices (HICP) is obtained from Eurostat, whereas food consumer prices are obtained from the OECD (Organisation for Economic Co-operation and Development), starting in 1971, and we link these with Eurostat data from 1996 onward.

- First, we construct a long time series of the energy component of the consumer price index for the European region by aggregating the energy indices of France, Germany, and Italy, which are available from the OECD database. This series is

²⁷For eggs, we use the Eurostat data until 1996.

Figure B.1. Comparison between Our Price Index and Eurostat Data



Note: The comparison starts from 2000 onward, since this is when Eurostat started reporting annual data. We show that our price index, which combines monthly data from Eurostat and the European Commission's DG AGRI, when aggregated at a yearly frequency, shows a very similar evolution to the original yearly price index of Eurostat.

considered a proxy for the energy index of the euro area. In the aggregation, we implement the Eurostat methodology.


- Second, we construct a long time series of the HICP energy component for the euro area by exploiting all available data. Data for the energy component of the consumer price index for the euro area are available from the OECD database from 1977 onward and from Eurostat from 1990 onward. To stretch the Eurostat series back to 1971, we proceed as follows. We extend the Eurostat energy index for the euro area backwards until 1977 by applying the growth rate of the OECD energy index for the euro area. Then, we extend the resulting series back to 1971 by applying the growth rate of the aggregate energy index based on OECD data for Italy, France, and Germany computed in the previous step.

- Then, once we have long time series for the (i) headline HICP, (ii) food HICP, and (iii) energy HICP for the euro area since 1971, we construct the core HICP index for the euro area by first adding up (ii) and (iii) and then subtracting the resulting sum from the headline HICP. Note, for these operations we implement the Eurostat methodology.

Other Control Variables. The following are additional control variables:

- **Crude Oil Prices and the Fertilizer Price Index.** These are obtained from the World Bank’s Pink Sheet and converted into euros.
- **Economic Activity for the Euro Area.** In the European region-based estimations, we also control for the economic activity in the European region by including as a regressor the monthly industrial production index (IPI) growth in the region. Since we need a large time series from the 1970s

Figure B.2. European and World Food Commodity Price Inflation



Note: Own calculations.

onward, we construct our own variable by combining different data. Most of the data are provided by the OECD, except for the most recent data (December 2023), which are obtained from the Eurostat database. In all cases, we consider industrial production excluding construction. For January 1971 to July 1974, we use the IPI for Europe (OECD), and we link this with the IPI for the euro area for August 1975 to November 2023. For the last data point (December 2023), which at the time of writing has not yet been released in the OECD database, we use industrial production excluding construction, obtained from Eurostat.

Table B.1. Components of the Global Food Commodity Price Index (WB)

Commodity Component	Time Availability
Rice	1960:M1–2023:M12
Wheat	1960:M1–2023:M12
Maize	1960:M1–2023:M12
Barley	1960:M1–2020:M8
Soybeans	1960:M1–2023:M12
Soybean Oil	1960:M1–2023:M12
Soybean Meal	1960:M1–2023:M12
Palm Oil	1960:M1–2023:M12
Coconut Oil	1960:M1–2023:M12
Groundnut Oil	1960:M1–2023:M12
Sugar	1960:M1–2023:M12
Bananas	1960:M1–2023:M12
Beef	1960:M1–2023:M12
Chicken	1960:M1–2023:M12
Oranges	1960:M1–2023:M12
Cocoa (*)	1960:M1–2023:M12
Coffee, Robusta (*)	1960:M1–2023:M12
Tea, Average Three Auctions (*)	1960:M1–2023:M12
Fish Meal (**)	1979:M1–2023:M12
Palm Kernel Oil (**)	1996:M1–2023:M12
Rapeseed Oil (**)	2002:M2–2023:M12
Sunflower Oil (**)	2002:M2–2023:M12
Lamb (**)	1971:M1–2023:M10
Mexican Shrimp (**)	1960:M1–2023:M10
Sorghum (***)	1960:M1–2020:M08
Groundnuts (***)	1980:M1–2023:M12

Note: These are the food commodity prices that compose the global food commodity price index as provided by the World Bank. (*) denotes subcomponents that are part of the World Bank Beverage Commodity Price Index. (**) denotes subcomponents that are not part of either the Food or the Beverage Commodity Index. (***) indicates weights considered in the Food Commodity Price Index.

Table C.1. Literature Review: Impact of El Niño on GDP Growth

	Results	Coverage	El Niño	Method	Sample	Controls
Beltran, Hadzi-Vaskov, and Stepanov (2023)	Heterogeneous effect on GDP across countries: negative for CAPDR(*), Colombia, Ecuador, Peru; positive for Argentina, Paraguay, Uruguay; ambiguous on Bolivia, Brazil, Chile, Mexico	Latin America	Oceanic Niño Index (ONI)	Local projections in cross-country panel	Not mentioned	Account for other shocks (but do not mention which ones)
Cashin, Mohaddes, and Raissi (2017)	Heterogeneous responses across countries: negative but short-lived for Australia, Chile, Indonesia, India, Japan, New Zealand, South Africa; positive for the U.S., Europe, China	Global: 21 countries/regions in an interlinked model	Southern Oscillation Index (SOI)	GVAR model	1979–2013	Inflation, real exchange rate, short-term and long-term interest rates, real equity prices, real energy and non-fuel commodity prices
Martín (2016)	GDP loss that ranges between 0.5 and 1.5 pp depending on the country, mostly concentrated in agriculture, infrastructure, and energy sector.	Bolivia, Peru, Ecuador, Colombia, Venezuela	Southern Oscillation Index (SOI)	Country-level AR models	Not defined, but includes 1982–98	Economic crisis dummies
Brunner (2002)	Positive effect of 0.5 pp on GDP growth in the G7 countries, significant at the 10 percent level. Possible explanations: increase in investment spending or in aggregate demand from other countries devastated by El Niño (e.g., food, housing materials, machinery, equipment).	Global: aggregate G7 countries	SST or SOI anomaly measure	VAR models	1963–98	Average CPI inflation rate, average GDP growth rate, international real commodity price inflation

(continued)

Table C.2. Literature Review: Impact of El Niño on Inflation

	Results	Coverage	Prices	El Niño	Method	Sample	Controls
Ubilava and Holt (2013)	Palm oil is directly affected by El Niño/La Niña (since it is produced in Oceania and Southeast Asia). Other vegetable oils (soybean oil, sunflower seed oil, rapeseed oil produced in not directly affected regions) are indirectly affected due to strong substitutability in consumption between vegetable oils. El Niño (La Niña) yields a vegetable oil price increase (decrease). A three-standard-deviation ENSO shock, 0.7 °C sea surface temperature (SST) anomaly, yields a 6–10 percent change in vegetable oil prices.	Global	Global vegetable oil prices: palm, soybean, rapeseed, and sunflower seed	SST anomalies	STAR	Monthly data, 1972–2010	Month dummies to control for seasonal effects
Cashin, Mohaddes, and Raissi (2017)	Non-fuel commodity prices increase after El Niño. Energy prices increase due to higher demand from Asia, to compensate for (i) lower electricity output and (ii) an increase in water demand for irrigation. This translates into a global oil demand shock (both oil prices and global output increase). CPI inflation increases in most countries (in the U.S. and Europe this increases by 0.12 and 0.07 pp, respectively). Responses are larger in economies that have a higher share of food in their CPI baskets.	21 countries/ regions (*)	National CPI inflation; global energy and non-fuel commodity prices	Southern Oscillation Index (SOI)	Country-specific local projections and GVAR	1979–2013	GDP growth, CPI inflation, exchange rate, short-term and long-term interest rates, equity prices, energy and non-fuel commodity prices

(continued)

Table C.2. (Continued)

	Results	Coverage	Prices	El Niño	Method	Sample	Controls
Martín (2016)	El Niño increases CPI inflation from 1 pp (Colombia) to 5 pp (Ecuador) depending on the country.	Bolivia, Peru, Ecuador, Colombia, and Venezuela	National CPI inflation	Southern Oscillation Index (SOI)	Country-level AR models	Includes the 1982–98 period; sample period not mentioned	Change in monetary aggregate as proxy for monetary policy
Romero and Saldarriaga (2024)	El Niño directly affects food inflation, which in turn increases inflation expectations (second round of effects).	Colombia	Inflation expectations, local food prices	Oceanic Niño Index (ONI)	BVARx	Monthly data, 2003–20	International food price index, real exchange rate, local food prices, industrial production, inflation expectations
Brunner (2002)	A one-standard-deviation positive surprise in ENSO raises real commodity price inflation by 3.5–4 pp. Qualitatively similar effects are observed on G7 overall prices, although much weaker and only marginally significant. This is because primary commodities account for only a fraction of overall finished good costs.	Global aggregate G7 countries	Global real non-oil primary commodity prices and G7 consumer price inflation	SST or SOI anomaly measure	VAR	Quarterly data, 1963–98	Average CPI inflation rate, average GDP growth rate, international real commodity price inflation


(continued)

Table C.2. (Continued)

	Results	Coverage	Prices	El Niño	Method	Sample	Controls
Abril-Salcedo, Melo-Velandia, and Parra-Amado (2020)	Cumulated effect of El Niño is a 600 bp increase in food inflation nine months after the shock. Over 1962–2018, no evidence emerges of changes in the size of food inflation responses over time. A La Niña shock yields ambiguous effects.	Colombia	National food price inflation	SST anomaly measure	Nonlinear smooth transition regression model (STR)	Monthly data, 1962–2018	No
Atems and Sardar (2021)	Heterogeneous effects between La Niña and El Niño. (**) La Niña shocks lead to positive responses to overall agriculture commodity prices. This is because agricultural output decreases after La Niña events, which increases prices in the following period. El Niño shocks do not affect overall agriculture commodity prices. Results for stocks are similar: La Niña shocks are beneficial to stock prices of agricultural companies, while El Niño shocks are detrimental or insignificant.	U.S.	Stock prices of 12 U.S. food and agricultural companies and global food commodity prices	SST anomaly measure	State-dependent local projections	Weekly/monthly data, 1990–2019	CPI price index, industrial production index


(continued)

Figure D.1. Response of International Commodity Prices to an Intensification of El Niño



Note: Monthly cumulative response of international food commodity prices to an El Niño shock. Full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations. The shock considered in the left (right) panel is an intensification of an El Niño, i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño event is already taking place.

Figure D.2. Response of International Commodity Prices to an Intensification of La Niña




Note: Monthly cumulative response of international food commodity prices to a La Niña shock. Full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations. The shock considered is an intensification of a La Niña, i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when a La Niña event is already taking place.

Appendix E. Disaggregated Effect of El Niño Shocks on EU Commodity and Consumer Prices


In this appendix, we analyze the disaggregated effects of El Niño shocks on EU commodity prices for the most recent period (1997–2023). While the aggregated impact is presented in Section 5.2 of the main text, here we delve into the disaggregated effects to demonstrate that the overall impact is disinflationary, despite being barely significant at the aggregate level. This disinflationary effect is primarily due to composition effects. In particular, meat and cereals are negatively and significantly affected by the intensification of El Niño, while dairy products are positively affected. We find no significant impacts on oils and fats, sugar, and coffee. Thus, the food aggregate containing the six food groups masks important impacts on the products with the greatest weight in the shopping basket. When we run the robustness exercise considering the entire value chain in Appendix F, these results become clearer. The negative impacts on meat and cereals become more pronounced, while the positive impact on dairy products diminishes and appears much later. Therefore, the disinflationary impact of El Niño is confirmed when taking into account the pass-through along the value chain.

Figure E.1. Disaggregated Response of EU Food Commodity Prices by Food Groups




Note: Monthly cumulative response of EU commodity prices according to our own six food group aggregate to an El Niño shock. Full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations. The shock considered is an intensification of an El Niño, i.e., an increase of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño event is already taking place.

Figure F.1. Response of Euro-Area Inflation: Value Chain (Eurostat aggregates)




Note: Monthly cumulative response of euro-area headline inflation and food inflation—Eurostat aggregate—to an El Niño shock (left panel) and a La Niña shock (right panel) along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1 °C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.

Figure F.2. Response of Euro-Area Food Inflation: Value Chain (own aggregate)



Note: Monthly cumulative response of euro-area food HICP according to our own six food group aggregate to an El Niño shock (left panel) and a La Niña shock (right panel) along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. In the third row, full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.


Figure F.3. Response of EU Food Commodity Prices: Value Chain



Note: Monthly cumulative response of European Union food commodity prices to an El Niño shock (left panel) and a La Niña shock (right panel), along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. In the third row, full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.


Appendix G. Robustness: Results for France, Italy, Germany, and Spain

Figure G.1. France: Inflation Responses




Note: Monthly cumulative response of France’s headline and food HICP to an El Niño/La Niña shock along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. Results are based on linear projection estimations. The shock considered is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.

Figure G.2. Germany: Inflation Responses




Note: Monthly cumulative response of Germany’s headline and food HICP to an El Niño/La Niña shock along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. Results are based on linear projection estimations. The shock considered is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.

Figure G.3. Italy: Inflation Responses



Note: Monthly cumulative response of Italy’s headline and food HICP to an El Niño/La Niña shock along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. Results are based on linear projection estimations. The shock considered is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.


Figure G.4. Spain: Inflation Responses



Note: Monthly cumulative response of Spain’s headline and food HICP to an El Niño/La Niña shock along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. Results are based on linear projection estimations. The shock considered is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.


Appendix H. Robustness: Historical Evidence

Figure H.1. Response of Euro-Area Headline and Food and Food Commodity Inflation: Data since 1970



Note: Monthly cumulative response of euro-area headline and food HICP, as well as food commodity prices, to an El Niño shock (left panel) and a La Niña shock (right panel) along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. Results are based on linear projection estimations. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.


Figure H.2. Response of Disaggregated European Union Food Commodity Prices: Data since 1970



Note: Monthly cumulative response of European Union food commodity prices to an El Niño shock (left panel) and a La Niña shock (right panel). Full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.


Appendix I. Robustness: Adding Controls and Lags

Figure I.1. Response of International Food Commodity Prices to Adding the Exchange Rate as a Control Variable



Note: Monthly cumulative response of international food commodity prices to an El Niño shock (left panel) and a La Niña shock (right panel) along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.


Figure I.2. Response of International Food Commodity Prices to Adding 12 Lags of Food Commodity, Fertilizer, and Oil Prices and Activity



Note: Monthly cumulative response of international food commodity prices to an El Niño shock (left panel) and a La Niña shock (right panel) along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place.


Appendix J. Robustness: Treating El Niño and La Niña Simultaneously

Figure J.1. Response of Euro-Area Headline Inflation:
Treating El Niño and La Niña Simultaneously




Note: Monthly cumulative response of euro-area headline HICP to an El Niño shock (left panel) and a La Niña shock (right panel) along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place. Both events are considered simultaneously, that is, we include both El Niño and La Niña variables in the model, as defined when we analyzed them individually.

Figure J.2. Response of Euro-Area Food Inflation:
Treating El Niño and La Niña Simultaneously




Note: Monthly cumulative response of euro-area food HICP to an El Niño shock (left panel) and a La Niña shock (right panel) along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. In the second row, full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place. Both events are considered simultaneously; that is, we include both El Niño and La Niña variables in the model, as defined when we analyzed them individually.

Figure J.3. Response of European Union Food Commodity Inflation: Treating El Niño and La Niña Simultaneously



Note: Monthly cumulative response of European Union food commodity to an El Niño shock (left panel) and a La Niña shock (right panel) along with 5 percent (lighter blue area) and 10 percent (darker blue area) confidence bands. In the second row, full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place. Both events are considered simultaneously, that is, we include both El Niño and La Niña variables in the model, as defined when we analyzed them individually.

Appendix K. Robustness: El Niño and La Niña Simultaneous and Complete
Figure K.1.1. Response of Disaggregated European Union Food Commodity
Prices: El Niño and La Niña Simultaneous and Complete



Note: Monthly cumulative response of European Union food commodity prices to an El Niño shock (left panel) and a La Niña shock (right panel). Full (empty) symbols indicate statistical significance at the 5 (10) percent level; a solid line indicates no statistical significance. Results are based on linear projection estimations. The shock considered in the left (right) panel is an intensification of an El Niño (La Niña), i.e., an increase (decrease) of 1°C in the ocean surface temperature in the east-central tropical Pacific when an El Niño (La Niña) event is already taking place. Note, in this specification the original ONI variables for El Niño and La Niña are included as such, which means that (i) both types of events are considered simultaneously, and (ii) all original values of the ONI variable are considered, and no censoring is applied.

Appendix L. EU Trade Patterns over Time

For the purpose of this exercise, we rely on annual data from the “Base pour l’Analyse du Commerce International of the Centre d’études prospectives et d’informations internationales” (BACI-CEPII), at the HS-6 product level. The data set includes bilateral trade flow information about the import value of each trade partner over a 28-year period from 1995 to 2022.

We aggregate import data into the six food categories that we used in the rest of the study: cereals, coffee, meat, dairy products and eggs, oils and fats, and sugar. In addition, the Economic and Monetary Union (EMU) is considered as a single importer, so the import value of each product coming from each trading partner is aggregated among the member states. Only imports from extra-EU economies are considered. We compute two measures.

First, we use the Herfindahl-Hirschman Index (HHI) to represent the degree of import concentration across trade partners. It ranges between 0 and 1, with highly concentrated imports resulting in a higher value of the indicator.²⁸ For each food category c , the HHI index is constructed by adding up the squared shares of EMU imports coming from country j in the total EMU imports of food category c :

$$HHI_{EMU}^c = \sum_{j=1}^N (Share_{EMU,j}^c)^2. \quad (L.1)$$

This general measure of concentration, HHI_{Total} , can be disaggregated to focus only on specific trading partners. In particular, the concentration of imports coming from economies affected by El Niño or La Niña, $HHI_{Ni\tilde{n}x}$, can be calculated by adding up the squared shares of EU imports coming only from countries affected by any of these climate phenomena. Similarly, the concentration of imports coming from the rest of the countries, $HHI_{no_Ni\tilde{n}x}$, can be calculated by adding up the squared import shares coming only from economies not affected by El Niño or La Niña.²⁹

²⁸An index equal to 1 represents a case of extreme concentration: a product that is imported from one trade partner only. An index close to 0 represents a case of high diversification: a product that is imported from many trade partners.

²⁹The sum of $HHI_{Ni\tilde{n}x}$ and $HHI_{no_Ni\tilde{n}x}$ equals HHI_{Total} .

The second measure we consider is the import share, which is computed for each food category by adding up the shares of EMU imports coming from a specific group of countries in the total EU imports of the corresponding food category. Thus, the index *Share_Niñx* can be calculated as the share of EMU imports coming from countries affected by El Niño or La Niña and the index *Share_no_Niñx* as the share of EMU imports coming from the rest of the economies.³⁰ Results for import concentration and import shares are shown in Figures 11 and 10 of Section 7 in the main text, respectively.

Appendix M. Construction of Newspaper-Based Index of ENSO Popularity

This appendix explains the construction of the newspaper-based indices discussed in Section 1 of the main text. These indices capture the importance ENSO phenomena have gained in the press over time in the following countries: U.S., Germany, France, Italy, and Spain.

To construct these indices, we build upon the methodology of Baker, Bloom, and Davis (2016), who constructed the well-known economic policy uncertainty (EPU) indicator for a number of countries.

To access newspaper articles, we rely on the Factiva repository provided by Dow Jones. We focus on the printed editions of newspapers and ignore their online versions. For each country, we consider as many newspapers as possible to ensure that selected newspapers cover almost the entire ideological spectrum and hence minimize any ideological bias that may stem from analyzing the media. We consider the following newspapers, by country:

- U.S., Canada, and United Kingdom: the *Boston Globe*, the *Telegraph*, the *Economist*, the *Globe and Mail*, the *Guardian*, the *Wall Street Journal*, *Los Angeles Times*, the *New York Times*, the *Times*, *Chicago Tribune*, and the *Washington Post*.

³⁰The sum of *Share_Niñx* and *Share_no_Niñx* is 1.

- Germany: *Berliner Zeitung*, *Berliner Morgenpost*, *Die Welt*, *Frankfurter Allgemeine Zeitung*, *Handelsblatt*, *Süddeutsche Zeitung*, *Der Tagesspiegel*, and *Die Tageszeitung*.
- France: *La Tribune*, *L'Opinion*, *Le Monde*, *Libération*, *La Croix*, *Le Figaro*, and *Les Echos*.
- Italy: *Corriere della Sera*, *La Stampa*, *Il Sole 24 Ore*, *ItaliaOggi*, *La Repubblica*, *Il Giornale*, and *Il Fatto Quotidiano*.
- Spain: *ABC*, *Cinco Días*, *El Economista*, *El Mundo*, *El País*, *Expansión*, and *La Vanguardia*.

For the U.S., Canada, and United Kingdom, the time coverage is longer such that we can construct the indicator from 1980. For the considered European countries (*the big 4*), our indicator starts from January 1997 due to data availability. We construct our measure of ENSO popularity at a monthly frequency and aggregate it at the yearly level for graphical purposes.

In a nutshell, the ENSO popularity index is constructed as the proportion of newspaper articles that are about ENSO climate phenomena. To do this, we count the number of articles in each month and newspaper that contain either *El Niño* or *La Niña* and which are also on the topic of the *climate*. For each country, we adapt the language accordingly. In practice, we consider the following keywords for each country:

- U.S., Canada, and United Kingdom: (*El Nin** or *La Nin**) and *climat**
- Germany: (*El Nin** or *La Nin**) and *klima**
- France: (*El Nin** or *La Nin**) and *clima**
- Italy: (*El Nin** or *La Nin**) and *clima**
- Spain: (*El Nin** or *La Nin**) and *climático** not (*Jesus* or *Navidad* or *Greta*). Note, since in Spanish *El Niño* means boy and *La Niña* means girl, we need to add additional requirements to minimize false positives (i.e., counting articles that are not related to the ENSO phenomena). In particular, we need to exclude articles about the climate and Christmas, which can contain *El Niño Jesus* (which means *baby Jesus*) and articles about climate and the young climate activist Greta Thunberg, which may contain *La Niña Greta*.

The construction of the indices closely follows the procedure used by Baker, Bloom, and Davis (2016). First, for each newspaper we count the number of articles published in a given month that contain the aforementioned keywords. Second, we divide this count by the total number of published articles that month, in order to express this count as a proportion of articles published by the newspaper in that month. Third, we standardize each monthly series of scaled counts by dividing by the standard deviation. This makes the volatility of the series comparable across newspapers. Fourth, for each country we average the newspaper-based standardized series across newspapers to compute an aggregated index. Fifth, we rescale the resulting index to mean 100 to obtain a set of homogeneous country-based ENSO popularity indices. Finally, these indicators are aggregated at a yearly frequency and plotted in Figure 2.

Note, the huge peak corresponding to the last El Niño event (2023–24) is striking, especially for the U.S. We double-checked the articles selected to construct the indicators for this period, and the peak is indeed genuine, i.e., all articles selected refer to El Niño or La Niña climate events when discussing topics related to climate. This confirms that the public is now extremely interested in this topic.

References

- Abril-Salcedo, D. S., L. F. Melo-Velandia, and D. Parra-Amado. 2020. “Nonlinear Relationship between the Weather Phenomenon El Niño and Colombian Food Prices.” *Australian Journal of Agricultural and Resource Economics* 64 (4): 1059–86. <https://doi.org/10.1111/1467-8489.12394>.
- Adams, R. M., L. L. Houston, B. A. McCarl, M. Tiscareño, J. Matus, and R. F. Weiher. 2003. “The Benefits to Mexican Agriculture of an El Niño–Southern Oscillation (ENSO) Early Warning System.” *Agricultural and Forest Meteorology* 115 (3-4): 183–94. [https://doi.org/10.1016/s0168-1923\(02\)00201-0](https://doi.org/10.1016/s0168-1923(02)00201-0).
- Adolfson, J. F., and M.-S. Lappe. 2023. “Risks to Global Food Commodity Prices from El Niño.” *Economic Bulletin* (European Central Bank) 6/2023: Box 1. <https://ideas.repec.org/a/ecb/ecbbox/202300061.html>.

- Andrian, L., C. M. Álvarez, A. Chávez, E. Díaz, C. Larrahondo, L. F. Serrudo, M. Alzamora, and D. Cárdenas. 2024. “Efectos del fenómeno “El Niño” en la Región Andina.” Technical Report, Interamerican Development Bank. <https://sinia.minam.gob.pe/sites/default/files/archivos/public/docs/Efectos-del-FEN-en-la-Region-Andina-una-aproximacion-empirica.pdf>.
- Atems, B., and N. Sardar. 2021. “Exploring Asymmetries in the Effects of El Niño–Southern Oscillation on U.S. Food and Agricultural Stock Prices.” *Quarterly Review of Economics and Finance* 81 (August): 1–14. <https://doi.org/10.1016/j.qref.2021.04.013>.
- Baker, S. R., N. Bloom, and S. J. Davis. 2016. “Measuring Economic Policy Uncertainty.” *Quarterly Journal of Economics* 131 (4): 1593–1636. <https://doi.org/10.1093/qje/qjw024>.
- Beltran, P., M. Hadzi-Vaskov, and I. Stepanov. 2023. “Online Annex 2. El Niño’s Potential Impact on Latin America.” In *Regional Economic Outlook—Western Hemisphere*, Online Annex 2. Washington, D.C.: International Monetary Fund. <https://www.imf.org/-/media/Files/Publications/REO/WH/2023/October/English/onlineannex-2-en.ashx>.
- Borrallo, F., L. Cuadro-Sáez, Á. Gras-Miralles, and J. J. Pérez. 2026. “The Transmission of Shocks to Food and Energy Commodity Prices to Food Inflation in the Euro Area.” *Applied Economics Letters* 33 (3): 411–16. <https://doi.org/10.1080/13504851.2024.2369711>.
- Brunner, A. D. 2002. “El Niño and World Primary Commodity Prices: Warm Water or Hot Air?” *Review of Economics and Statistics* 84 (1): 176–83. <http://www.jstor.org/stable/3211747>.
- Cashin, P., K. Mohaddes, and M. Raissi. 2017. “Fair Weather or Foul? The Macroeconomic Effects of El Niño.” *Journal of International Economics* 106 (May): 37–54. <https://doi.org/10.1016/j.jinteco.2017.01.010>.
- Cheng, I.-H., and W. Xiong. 2014. “Financialization of Commodity Markets.” *Annual Review of Financial Economics* 6 (1): 419–41. <https://doi.org/10.1146/annurev-financial-110613-034432>.

- Cuadro-Sáez, L., J. J. Pérez, and F. Viani. 2025. "Food Commodity Production and Imports in Spain and the EU: How Much External Autonomy is There?" *Economic Bulletin* (Bank of Spain) 2025/Q2: Article 04. <https://doi.org/10.53479/39605>.
- European Central Bank. 2011. "Financialisation of Commodities." Box 4 in Financial Stability Review (December). https://www.ecb.europa.eu/press/financial-stability-publications/fsr/foqus/2011/pdf/ecb~6fdfdfce1c.fsrbox201112_04.pdf.
- . 2023. "Eurosystem Staff Macroeconomic Projections for the Euro Area." Technical Report. https://www.ecb.europa.eu/press/projections/html/ecb.projections202312_eurosystemstaff~9a39ab5088.en.html.
- Ferreira, G. L. M., J. M. Tonin, and A. F. Alves. 2022. "Impacts of El Niño Southern Oscillation on Hedge Strategies for Brazilian Corn and Soybean Futures Contracts." *Revista de Economia e Sociologia Rural* 60 (spe). <https://doi.org/10.1590/1806-9479.2021.250643>.
- Ferrucci, G., R. Jiménez-Rodríguez, and L. Onorante. 2012. "Food Price Pass-Through in the Euro Area: Non-Linearities and the Role of the Common Agricultural Policy." *International Journal of Central Banking* 8 (1, March): 179–218. <https://www.ijcb.org/journal/ijcb12q1a9.htm>.
- Food and Agriculture Organization (FAO). 2023. "El Niño: Anticipatory Action and Response Plan, October 2023–March 2024." Technical Report, Food and Agriculture Organization of the United Nations. <https://doi.org/10.4060/cc8496en>.
- Garzon, I. 2006. "The MacSharry Reform (1992)." In *Reforming the Common Agricultural Policy*, 61–75. Palgrave Macmillan UK. https://doi.org/10.1057/9780230626577_6.
- Hsiang, S. M., and K. C. Meng. 2015. "Tropical Economics." *American Economic Review* 105 (5): 257–61. <https://doi.org/10.1257/aer.p20151030>.
- Iizumi, T., J.-J. Luo, A. J. Challinor, G. Sakurai, M. Yokozawa, H. Sakuma, M. E. Brown, and T. Yamagata. 2014. "Impacts of El Niño Southern Oscillation on the Global Yields of Major Crops." *Nature Communications* 5 (1): 3712. <https://doi.org/10.1038/ncomms4712>.
- Jacks, D. S., K. H. O'Rourke, and J. G. Williamson. 2009. "Commodity Price Volatility and World Market Integration since

- 1700.” NBER Working Paper No. 14748. https://www.nber.org/system/files/working_papers/w14748/w14748.pdf.
- Jiang, J., and T. R. Fortenbery. 2019. “El Niño and La Niña Induced Volatility Spillover Effects in the U.S. Soybean and Water Equity Markets.” *Applied Economics* 51 (11): 1133–50. <https://doi.org/10.1080/00036846.2018.1524980>.
- Jordà, Ò. 2005. “Estimation and Inference of Impulse Responses by Local Projections.” *American Economic Review* 95 (1): 161–82. <https://doi.org/10.1257/0002828053828518>.
- Kreiss, J.-P., and S. N. Lahiri. 2012. “Bootstrap Methods for Time Series.” *Handbook of Statistics*, Vol. 30, ed. T. S. Rao, S. S. Rao, and C. R. Rao, 3–26. Elsevier. <https://doi.org/10.1016/B978-0-444-53858-1.00001-6>.
- Laosuthi, T., and D. D. Selover. 2007. “Does El Niño Affect Business Cycles?” *Eastern Economic Journal* 33 (1): 21–42. <https://doi.org/10.1057/ej.2007.2>.
- Martín, L. 2016. “¡Es Niño!: Impacto económico en la Región Andina.” <https://doi.org/10.18235/0000234>.
- Montiel Olea, J. L., and M. Plagborg-Møller. 2021. “Local Projection Inference is Simpler and More Robust than You Think.” *Econometrica* 89 (4): 1789–1823. <https://arxiv.org/pdf/2007.13888>.
- Pantoja-Robayo, J., and D. Rodríguez-Guevara. 2023. “The Climate Effect on Colombian Coffee Prices and Quantities Based on Risk Analysis and the Hedging Strategy in Discrete Setting Approach.” *Agris on-line Papers in Economics and Informatics* 15 (4): 97–107. <https://doi.org/10.7160/aol.2023.150407>.
- Paz, J. O., C. W. Fraisse, L. U. Hatch, A. García y García, L. C. Guerra, O. Uryasev, J. G. Bellow, J. W. Jones, and G. Hoogenboom. 2007. “Development of an ENSO-based Irrigation Decision Support Tool for Peanut Production in the Southeastern US.” *Computers and Electronics in Agriculture* 55 (1): 28–35. <https://doi.org/10.1016/j.compag.2006.11.003>.
- Peersman, G. 2022. “International Food Commodity Prices and Missing (Dis)Inflation in the Euro Area.” *Review of Economics and Statistics* 104 (1): 85–100. https://doi.org/10.1162/rest_a_00939.
- Raj, E. E., R. R. Kumar, and K. V. Ramesh. 2020. “El Niño–Southern Oscillation (ENSO) Impact on Tea Production and

- Rainfall in South India.” *Journal of Applied Meteorology and Climatology* 59 (4): 651–64. <https://doi.org/10.1175/jamc-d-19-0065.1>.
- Romero, J. V., and S. N. Saldarriaga. 2024. “Weather Shocks and Inflation Expectations in Semi-Structural Models.” *Latin American Journal of Central Banking* 5 (2): Article 100112. <https://doi.org/10.1016/j.latcb.2023.100112>.
- Smith, S. C., and D. Ubilava. 2017. “The El Niño Southern Oscillation and Economic Growth in the Developing World.” *Global Environmental Change* 45 (July): 151–64. <https://doi.org/10.1016/j.gloenvcha.2017.05.007>.
- Tang, K., and W. Xiong. 2012. “Index Investment and the Financialization of Commodities.” *Financial Analysts Journal* 68 (6): 54–74. <https://doi.org/10.2469/faj.v68.n6.5>.
- Ubilava, D. 2011. “El Niño, La Niña, and World Coffee Price Dynamics.” *Agricultural Economics* 43 (1): 17–26. <https://doi.org/10.1111/j.1574-0862.2011.00562.x>.
- . 2017. “The ENSO Effect and Asymmetries in Wheat Price Dynamics.” *World Development* 96 (August): 490–502. <https://doi.org/10.1016/j.worlddev.2017.03.031>.
- . 2023. “What This Year’s El Niño Means for Wheat and Global Food Supply.” *The Conversation*. <https://theconversation.com/what-this-years-el-nino-means-for-wheat-and-global-food-supply-209386>.
- Ubilava, D., and M. Holt. 2013. “El Niño Southern Oscillation and Its Effects on World Vegetable Oil Prices: Assessing Asymmetries Using Smooth Transition Models.” *Australian Journal of Agricultural and Resource Economics* 57 (2): 273–97. <https://doi.org/10.1111/j.1467-8489.2012.00616.x>.
- Yeager, S. G., N. Rosenbloom, A. A. Glanville, X. Wu, I. Simpson, H. Li, M. J. Molina, et al. 2022. “The Seasonal-to-Multiyear Large Ensemble (SMYLE) Prediction System Using the Community Earth System Model Version 2.” *Geoscientific Model Development* 15 (16): 6451–93. <https://doi.org/10.5194/gmd-15-6451-2022>.

Zebiak, S. E., B. Orlove, A. G. Muñoz, C. Vaughan, J. Hansen, T. Troy, M. C. Thomson, A. Lustig, and S. Garvin. 2015. “Investigating El Niño-Southern Oscillation and Society Relationships.” *WIREs Climate Change* 6 (1): 17–34. <https://doi.org/10.1002/wcc.294>.

Low for Long: Could Extended Periods of Ultra-low Monetary Policy Rates Have Harmful Effects?*

Etibar Jafarov^a and Enrico Minnella^b

^aInternational Monetary Fund

^bBank of England

Extended periods of ultra-low monetary policy rates in advanced economies have rekindled debates about the zombification of weak companies and its impact on resource allocation, economic growth, inflation, and financial stability. Using both firm-level and macroeconomic data, we find that recessions are a critical factor in the increase in the number of zombie firms and that extended periods of low interest rates are associated with a higher probability of zombification. Our findings imply a trade-off between conducting a countercyclical monetary policy and using expansionary tools for long periods, which may lead to a combination of low interest rates, low growth, and high financial vulnerability. We also find that strong banking supervision reduces the probability of zombification.

JEL Codes: E44, E52, E58, G20.

1. Introduction

During the Global Financial Crisis (GFC) in 2008 and the COVID-19 pandemic in early 2020, central banks in advanced economies

*The authors would like to thank Mauricio Villafuerte and Gustavo Adler for supporting this research project. We would also like to thank Andrea Brasili, Fuad Hasanov, Fumitaka Nakamura, Beniamino Pisicoli, Irene Rizzoli, and Tammaro Terracciano for their helpful comments and suggestions, as well as participants at several IMF seminars and the 2023 CEMLA conference for their valuable comments. Finally, the authors would like to thank Mr. Nicolas Landeta and Ms. Daria Kolpakova for their excellent organizational and research support. Any views expressed are solely those of the authors and so cannot be taken to represent those of the International Monetary Fund, its Executive Board, or IMF management, the Bank of England or any of its committees, or to state their policies. All remaining errors are ours.

lowered policy rates effectively to the zero lower bound (ZLB) and implemented ultra-low interest rate policies alongside unconventional monetary policy (UMP) operations.¹ These policies are widely considered successful in preventing financial meltdowns and closing output gaps, as highlighted by Vayanos and Vila (2009), Krishnamurthy and Vissing-Jorgensen (2011), Harrison (2012, 2017), and Corrado, Grassi, and Minnella (2021), among others. However, in some jurisdictions, they were subsequently continued to boost aggregate demand amid subdued inflationary pressure. Critics suggest that such prolonged ultra-low monetary policy rates could lead to unsustainable capital flows and booms in housing and financial asset prices, raising the odds of a future financial crisis (Turner et al. 2019), supporting weak companies at the cost of healthy ones, and leading to financial fragility and resource misallocation, as discussed in frameworks such as the Austrian business cycle theory and Minsky's financial instability hypothesis (Minsky 1977). Japan's experience with the short-term policy rate at the ZLB over many years illustrates the risks associated with such policies. Attempts to combat deflation and low growth with various economic stimuli, running large fiscal deficits and ultra-loose monetary policy, led to the so-called zombie lending, as coined by Caballero, Hoshi, and Kashyap (2008). This phenomenon laid the foundation for a rich stream of literature on the causes and effects of corporate zombification. Their seminal paper estimates that around 30 percent of Japanese firms were recipients of zombie lending in the late 1990s.² Other examples in this wide literature include Acharya et al. (2019), who found that the Outright Monetary Transactions programs (OMT) in the euro zone during the Sovereign Debt Crisis benefited low-quality borrowers, contributing to the zombification of weak firms and negative spillovers to more productive firms; Banerjee and Hofmann (2018), who suggested that the increase in zombie firms may be linked to reduced financial pressure; and Bonfim et al. (2023), who found that

¹See Joyce et al. (2012) for a detailed review of the theoretical transmission channels of UMPs.

²Peek and Rosengren (2005) argue that zombie lending was not only a by-product of economic conditions but also contributed to Japan's low output growth during its "lost decade."

inspected banks are less likely to refinance zombie firms after the inspection.³

In this paper, we analyze whether extended periods of low monetary policy interest rates in advanced economies have contributed to the zombification of weak firms and low economic growth. Using both firm-level and macroeconomic data, we find that recessions are a critical factor in the increase in the number of zombie firms and that extended periods of low interest rates are associated with a higher probability of zombification. We also find that, in such environments, small and medium enterprises are more likely to become zombie firms since they are less dynamic than other firms and are usually in sectors least exposed to foreign competition. Moreover, we find that Japan experienced a larger economic slowdown than other countries during our sample period of 1996–2021, partly due to the increase in the number of zombie firms. This raises concerns about the sustainability of prolonged periods of ultra-low monetary policy rates, especially in countries with low growth rates. Additionally, we find evidence of negative spillovers to more productive companies operating in sectors with higher shares of zombie firms. Lastly, we find that higher supervisory power can play a role in reducing zombie lending, highlighting the importance of the post-GFC prudential regulations.

This paper makes at least two contributions to the literature. First, it uses both macroeconomic and large microeconomic data to study zombification trends and their impact on resource allocation and growth. Second, it analyzes the impact of the “low for long” policies on zombification. The paper is organized as follows: Section 2 describes the data and methods used and presents the results of various empirical analyses of factors of zombification, including extended periods of low monetary policy rates. Section 3 discusses the impact of zombification on economic growth. Section 4 discusses the impact of supervisory power on zombification and reflects on financial stability aspects. Section 5 concludes.

2. Empirical Analysis of Zombification

Answering questions related to zombie firms has one initial challenge related to defining such firms. Although zombies can certainly be

³See Acharya et al. (2022) for a more detailed literature review on the topic.

considered unprofitable corporates, Caballero, Hoshi, and Kashyap (2008) depart from this assumption and classify firms as zombies based on whether they have received subsidized credit. They do so by calculating a hypothetical lower bound on interest payments that they expect from high-quality borrowers and comparing it with the actual interest payment. Giannetti and Simonov (2013) and Acharya et al. (2019) use a similar definition. On the other hand, Banerjee and Hofmann (2018) focus on the profitability of firms by considering two different definitions. A broader measure, following Adalet McGowan, Andrews, and Millot (2018), considers firms as zombies if their interest coverage ratio (ICR)—defined as the ratio between earnings before interest and taxes (EBIT) and interest payments—is less than unity for at least 3 consecutive years and if they are at least 10 years old, and a narrower measure if the firm's Tobin's q is lower than its sector median. While the former focuses on the firm's current profitability (more precisely, on the firm's ability to pay its interest expenses), the latter focuses on future profitability. For our analysis, we follow the second strand of literature and inspect firms' profitability. We use firm-level annual data for publicly traded companies and small firms from the Orbis database, provided by Bureau Van Dijk, with around half a million firm-year observations.⁴ The data set spans 1996 to 2021 and contains balance sheet data and ratios for around 60,000 companies from 21 countries, including the United States, Japan, and all 19 euro-zone countries.⁵ According to our definition, a firm is classified as a zombie if it is at least 10 years old and its ICR is lower than one for at least two consecutive periods. The first requirement allows us to exclude relatively young, healthy corporates which might suffer initial losses due to their high investments—some of which through high loan volumes—in their first years of activity. In addition, we allow firms to fluctuate between the zombie and non-zombie states whenever the conditions above are not met. By doing so, we can detect the year-specific characteristics

⁴We merge firm-level data and information coming from different Orbis sub-databases and clean the data set following the steps outlined in Diez, Fan, and Villegas-Sánchez (2019).

⁵Data availability is very limited for Slovenian firms compared to the rest of the sample.


of a firm and link them to the current monetary policy implementations to inspect their correlation in recession and ZLB periods. To check the robustness of our results, we also repeat the analyses using an alternative definition of zombie firms, under which a company's ICR is below one in three successive years, which produces similar results (see the appendix for details).

2.1 Stylized Facts

As shown in Figure 1, the number of zombie firms and their share in the total number of firms increased after recessions and broadly declined at other times. In the U.S., the share of zombie firms increased during the recessions in the early 2000s (after the collapse of the speculative dot-com bubble and the September 11 terrorist attacks), during the GFC, and during the COVID-19 pandemic, when the output gap turned negative. The share gradually declined in most other periods. In Japan, the share of zombie firms had a similar pattern, but Japan had high levels of zombie share also before our sample period, as documented by Caballero, Hoshi, and Kashyap (2008). In Europe, the zombie share increased to the highest level during the GFC and increased the most (among the three regions) during the pandemic. On the other hand, monetary policy rates, here proxied by the shadow rates,⁶ were reduced following the recessions and were tightened as the economies recovered. These episodes imply negative correlations between the shadow rate and zombie share, but there were also periods, such as 2011–13 in the

⁶The shadow rate, first suggested in Black (1995)'s seminal work, corresponds to the shortest-maturity interest rate that the yield curve would have generated had the ZLB not been binding. The literature offers many alternatives to estimating this measure. For example, Krippner (2013), Wu and Xia (2016), and Lemke and Vladu (2017) employ term structure models, while Lombardi and Zhu (2018) use factor analysis, and Mavroeidis (2021) and Ikeda, Mavroeidis, and Zanetti (2024) develop a new vector autoregression (VAR) estimation technique which nests the estimation of a shadow rate. Many theoretical and empirical papers have employed such instruments in the recent past. Among others, Mouabbi and Sahuc (2019) and Hohberger, Priftis, and Vogel (2019) estimate dynamic stochastic general equilibrium (DSGE) models, with the first using a shadow EONIA (euro overnight index average) rate and the second with an occasionally binding constraint on the ZLB by employing a shadow interest rate. They both build macroeconomic scenarios with low nominal interest rates and assess the effect of UMPs.


Figure 1. Share of Zombie Firms and the Shadow Rate



Note: The zombie shares (left-hand-side axis) are computed as the fraction of zombie firms over the total number of firms at time t . The shadow rates (right-hand-side axis) are Krippner (2013)’s estimates.
Sources: Orbis database and the authors’ calculations.

U.S., 2011–15 in Japan, and 2012–15 in the euro area, when both variables moved in the same direction. Note that in Japan, the shadow rate has been in the negative area since the early 2000s,


Figure 2. Zombie vs. Non-zombie Firms Performance



Note: The solid blue line is the mean value of the corresponding performance indicator for the firms classified as zombies, whereas the dashed red line is the mean value of the firms not classified as zombies in our data set.
Sources: Orbis database and the authors’ calculations.

except for a brief period in 2007–08. In the euro zone, the shadow rate turned negative during the GFC and stayed at negative levels for most years after that, when the share of zombie firms also declined. These point to the importance of controlling for the business cycle, which we do by including the recession variable, GDP growth rate, and output gap in various specifications. As expected, the data set reveals that zombie firms’ financial performance is worse than that of non-zombie firms. Almost by definition, zombie firms are more indebted, and most make losses. They also record lower sales and employment growth (Figure 2). But it is worth noting that most of them were still increasing their sales, employment, and borrowing except for the aftermath of the GFC, suggesting that they might have taken market shares from their healthier competitors.

Figure 3. Zombie Shares in Small-Medium and Large Enterprises




Note: The solid blue line and the dashed red line correspond to the zombie shares of SMEs and LEs, respectively.
Sources: Orbis database, and authors’ calculations.

Descriptive statistics from our data set also reveal an interesting pattern for smaller companies compared to larger enterprises. We identify the former as firms with 1,000 or fewer employees. Figure 3 shows the ratios of identified zombie firms over the total number of firms for both small and medium enterprises (SMEs) and large enterprises (LEs). In all regions, both series closely follow the path of the economic cycle, increasing during recessions and decreasing during expansions, and depict a clear outperformance of the LEs, having lower zombie shares throughout the sample period. However, there are some notable differences across regions. First, the U.S. experiences a larger gap between SMEs and LEs, with the latter being on average 10 percentage points (pp) less than the former. Second, the gap spikes in Japan around 2000, coinciding with the beginning of the Zero Interest Rate Policy (ZIRP) implemented by the Bank of Japan, and persists until the end of the sample period, although somewhat narrowing. Third, a similar gap emerges in the euro area in the early 2000s, increasing even further during the GFC. The gap narrows around 2020 but then increases again during the COVID-19 pandemic. On average, SMEs have a higher share of zombie firms.

Lastly, we are interested in the degree of persistence that the firms have in remaining zombie firms, once classified as such. Figure 4 illustrates the share of firms that exited the zombie state over the

Figure 4. Exiting Rate from the Zombie State



Note: The solid lines are the exiting rate for the three distinctive regions of the sample, and the dashed line is the full-sample exiting rate.
Sources: Orbis database, and authors’ calculations.

total number of zombie firms in each year and regional unit of our sample, as well as the share for the full sample. We compute it as the ratio of the firms not classified as zombies in t , while being classified as such in $t - 1$, over the total number of zombie firms in $t - 1$. To a certain extent, it can be observed how the capacity of a firm to exit the zombie state has some degree of correlation with the state of the economy. All regions experienced a massive drop between 2008 and 2010, as well as around 2013. This means that during crisis periods the number of firms already classified as zombies that are able to escape that state is significantly lower than average. Despite a similar pattern with the business cycle, Japan’s values are more volatile, with a higher average exit rate but a similar trough during the GFC.

2.2 Monetary Policy and Zombification

In our first regression analysis, we investigate the effect of monetary policy implementations on the probability of a firm becoming

a zombie. Our baseline specification for firm i in country j at time t reads as

$$\begin{aligned} \text{Zombie}_{i,j,t} = & \beta_1 \cdot \text{Shadow}_{j,t} + \beta_2 \cdot \text{Rec}_{j,t} + \beta_3 \cdot \text{ZLB}_{j,t} \\ & + \beta_4 \cdot \text{Shadow}_{j,t} \cdot \text{Rec}_{j,t} + \beta_5 \cdot \text{Shadow}_{j,t} \cdot \text{ZLB}_{j,t} \quad (1) \\ & + \beta_6 \cdot \text{SME}_{i,j,t} + \gamma \cdot X_{i,j,t} + u_{i,j,t}, \end{aligned}$$

where $\text{Zombie}_{i,j,t}$ is a dummy variable equal to one if a firm is classified as a zombie according to our definition, $\text{Shadow}_{j,t}$ is the shadow interest rate estimated using term structure models as estimated by Krippner (2013),⁷ and $\text{Rec}_{j,t}$ and $\text{ZLB}_{j,t}$ are dummy variables that are equal to one in recession periods and when the interest rate is lower than 25 basis points (bps), respectively. We also include their respective interaction terms with the shadow rate. We include an $\text{SME}_{i,j,t}$ dummy variable for small and medium enterprises, and a set of country-level and firm-level control variables $X_{i,j,t}$, including real GDP growth, the one-year lag of the shadow rate, return on firms' assets, and their profit margin (for each firm). Table 1 shows the results of our first estimation exercise. The first four columns illustrate panel fixed-effect (FE) estimates according to four different specifications. The first regression includes the recession and its interaction term with the shadow rate, while the second regression includes the ZLB dummy variable and its interaction term with the shadow rate. The third column includes both dummies and their interaction terms, and the fourth column augments this last specification with a lag of the shadow rate. In all settings, we control for firm size, profitability, balance sheet characteristics, and economic growth measured by GDP.

The estimation highlights a positive relationship between our dependent variable and recessions, and a negative relationship with the shadow rate when the policy rate is above the ZLB. This means that being in a recession increases the probability of a firm being a zombie firm. Being a small or medium enterprise also increases this probability.⁸ On the contrary, increases in aggregate output as

⁷We use the shadow rates as estimated by Krippner (2013) since they are the most comprehensive coverage of the countries/regions studied here.

⁸This outcome might be because smaller firms often belong to the sectors least exposed to foreign competition and thus may be less productive and more

Table 1. Effect on the Probability of Being a Zombie Firm during Recessions and at the ZLB

Dependent Variable: Zombie	Linear Probability Model				Probit Model (<i>mfx</i>)	
	(1)	(2)	(3)	(4)	(5)	(6)
Shadow Rate	-0.25*** (0.03)	-0.47*** (0.05)	-0.52*** (0.04)	-0.51*** (0.05)	-0.25*** (0.02)	-0.25*** (0.03)
Recession	1.19*** (0.10)		0.83*** (0.09)	0.54*** (0.12)	0.41*** (0.06)	0.08 (0.07)
Shadow Rate $\times$ Recession	0.24*** (0.04)		0.30*** (0.04)	0.23*** (0.04)	0.13*** (0.02)	0.07*** (0.03)
ZLB		1.06*** (0.16)	0.95*** (0.17)	1.10*** (0.19)	0.45*** (0.09)	0.74*** (0.11)
Shadow Rate $\times$ ZLB		1.17*** (0.82)	1.07*** (0.08)	1.17*** (0.08)	0.69*** (0.04)	0.75*** (0.05)
SME	0.59** (0.28)	0.53* (0.28)	0.55* (0.28)	0.51 (0.31)	0.31*** (0.09)	0.65*** (0.10)
GDP Growth	-2.40** (1.18)	-5.90*** (1.00)	-3.49*** (1.16)	-2.88** (1.23)	-1.44** (0.60)	-1.85*** (0.66)
ROA	-0.47*** (0.04)	-0.47*** (0.01)	-0.47*** (0.01)	-0.52*** (0.01)	-0.43*** (0.01)	-0.50*** (0.01)
Profit Margin	-0.22*** (0.01)	-0.22*** (0.01)	-0.21*** (0.01)	-0.21*** (0.01)	-0.06*** (0.01)	-0.05*** (0.01)
Shadow Rate (-1)				0.10*** (0.04)		0.14*** (0.03)
Firm FE	Yes	Yes	Yes	Yes	No	No
Observations	433,151	433,151	433,151	371,626	433,151	371,626
Individual Firms	63,488	63,488	63,488	54,813	63,488	54,813
R^2	0.09	0.09	0.09	0.10		

Note: Coefficients expressed in percentage points. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

well as the ROA (return on assets) and the profit margin of a firm, which control for firm-level health, reduce the probability of being a zombie firm. The impact of monetary policy decisions, represented by the shadow rate, on zombification depends on the economic cycle when these decisions are made. In particular, while the shadow rate is negatively correlated with the probability of a firm being a zombie under the first specification, this correlation turns positive when we control for ZLB environments under the second specification. Specifically, the interaction of the shadow rate with the ZLB variable (in the second specification) has a positive coefficient, which exceeds the negative coefficient for the shadow rate. Hence the relationship between the shadow rate and the dependent variable is positive. In particular, in this specification, a 1 pp decrease in the shadow rate seems to increase by roughly 0.5 pp the probability of a firm becoming a zombie, but at the ZLB, *ceteris paribus*, this probability decreases by around 0.6 pp. In other words, an expansionary monetary policy implementation at the ZLB reduces the probability of a firm being a zombie, but such a policy would lead to more zombification in the other (non-ZLB) periods. The third and fourth specifications, which include all the variables and the lag of the shadow rate variable, yield similar results. Columns 5 and 6 show the results with the same set of variables using a panel probit model. These results are consistent with the linear probability model (LPM) estimates (in the previous columns) both in terms of the sign and magnitude.

To check the robustness of our results, we also analyze the relationship between monetary policy implementation and the rise of zombie firms in a macroeconomic setting. First, we build the zombie share variable, a country-level measure of the number of zombies in country j at time t over the total number of firms. Then, we regress it against a set of relevant macroeconomic variables while allowing for country-specific fixed effects. Table 2 shows that, on average, being in a recession or at the ZLB increases the zombie share. However, when expansionary measures are implemented when the economy reaches the ZLB, the effects are beneficial to those firms, confirming the results of our previous estimation exercise.

exposed to domestic economic cycles. The reasons for this finding are not pursued further here.

Table 2. Effect of Macroeconomic Variables on the Country-Specific Zombie Shares

Dependent Variable: Zombie Share	(1)	(2)	(3)
Shadow Rate	0.25*** (0.04)	−0.33*** (0.08)	−0.36*** (0.07)
Rec	1.97*** (0.21)		1.65*** (0.20)
Shadow Rate × Rec	0.34*** (0.09)		0.34*** (0.08)
ZLB		1.05*** (0.36)	0.40 (0.30)
Shadow Rate × ZLB		2.00*** (0.16)	1.73*** (0.15)
GDP Growth	−3.99** (1.70)	−6.85*** (1.50)	−1.21 (1.58)
Inflation	0.24** (0.11)	0.33*** (0.26)	0.17* (0.09)
Country FE	Yes	Yes	Yes
Observations	546	546	546
Year-Observations	21	21	21
R^2	0.32	0.37	0.46
Note: Coefficients expressed in percentage points. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.			

2.3 Monetary Policy: Low for Long

While the previous subsection found that expansionary policy is generally associated with higher zombification but can help reduce zombification at the ZLB, this section aims to assess the impact of prolonged periods of low interest rates. We do so by introducing a variable named $Long_{j,t}$, a dummy variable for country j at time t , capturing periods when the short-term policy rate was kept low for an extended period. We set it equal to one if two conditions apply. Firstly, if the policy rate was lower than 25 bps for at least two consecutive years. Secondly, since unconventional monetary policies can on their own exert some effect on the probability of a firm becoming

a zombie, we decide to set the dummy equal to one at time t if also quantitative easing (QE) was not implemented. This second requirement ensures that we can capture the implementation of persistently low interest rate policies⁹ without mistakenly including the effects of QE on firms' borrowing costs and other financial indicators.¹⁰ For this analysis, we estimate the following regression:

$$\begin{aligned} \text{Zombie}_{i,j,t} = & \beta_1 \cdot \text{Shadow}_{j,t} + \beta_2 \cdot \text{Long}_{j,t} + \beta_3 \cdot \text{Shadow}_{j,t} \cdot \text{Long}_{j,t} \\ & + \beta_4 \cdot \text{Rec}_{j,t} + \beta_5 \cdot \text{SME}_{i,j,t} + \gamma \cdot X_{i,j,t} + u_{i,j,t}, \end{aligned} \quad (2)$$

where $X_{i,j,t}$ includes the same set of control variables as in our baseline specification. Table 3 reports the results of panel ordinary least squares (OLS) and probit model estimations. Both estimations confirm the negative relationship between the shadow rate and the probability of being a zombie firm. Long periods of low interest rates seem to be correlated with higher probabilities of having zombie firms, given that the interaction term of the monetary policy rate during prolonged periods of low interest rates is highly negative and significant.

These results highlight the negative impact that prolonged periods of very low interest rates can have on the rise of zombie firms. As discussed by Caballero, Hoshi, and Kashyap (2008), when interest rates are very low and banks are subject to strict capital requirements, their willingness to keep lending to unprofitable firms can generally increase, as they may hope for government support to the borrowers. Furthermore, expectations play an important role. The longer interest rates are kept low, the more likely the agents will believe that the central bank will stick to a low interest rate policy,

⁹Note that the purpose here is not to identify monetary policy shocks. Instead, we focus on the correlation between extended periods of low interest rates and zombification.

¹⁰In addition, asset purchases of central banks normally focus on buying financial instruments of healthier large firms experiencing some liquidity issues. Accordingly, QE might not be detrimental to productivity, while lower interest rates for extended periods would facilitate the survival of zombie firms, increasing the probability of a weak firm becoming a zombie. There are also alternative views. For example, Krishnamurthy and Vissing-Jorgensen (2011) highlight that the default risk channel of QE could benefit riskier firms more than healthier firms, increasing the probability of their survival and thus the probability of becoming zombie firms.

Table 3. Effect on the Probability of Being a Zombie Firm during Long Periods of Low Interest Rates

Dependent Variable: Zombie	Linear Probability Model		Probit Model (<i>mfx</i>)	
	(1)	(2)	(3)	(4)
Shadow Rate	−0.21*** (0.03)	−0.22*** (0.03)	−0.03*** (0.01)	−0.04*** (0.01)
Long	2.49*** (0.32)	1.97*** (0.32)	1.24*** (0.14)	0.94*** (0.14)
Shadow Rate × Long	−3.55*** (0.45)	−3.04*** (0.44)	−1.70*** (0.24)	−1.43*** (0.24)
Recession		1.02*** (0.10)		0.55*** (0.06)
SME	0.57** (0.28)	0.58** (0.28)	0.34*** (0.09)	0.34*** (0.09)
GDP Growth	−8.78** (0.98)	−3.75** (1.10)	−4.48*** (0.54)	−1.68*** (0.61)
ROA	−0.47*** (0.01)	−0.47*** (0.01)	−0.43*** (0.01)	−0.43*** (0.01)
Profit Margin	−0.22*** (0.01)	−0.22*** (0.01)	−0.06*** (0.01)	−0.06*** (0.01)
Firm FE	Yes	Yes	No	No
Observations	433,151	433,151	433,151	433,151
Individual Firms	63,488	63,488	63,488	63,488
R^2	0.09	0.09		
Note: Coefficients expressed in percentage points. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.				

even exacerbating this mechanism. This is mainly captured by the Japanese experience; Japan was the first country to adopt ultra-low monetary policy rates for extended periods, with the implementation of the ZIRP before extensively using QE—the first such case in an advanced economy.

To show the robustness of our results, we saturate our previous specification by controlling for all cross-sectional and time variations within firm, sector, and country. Table 4 shows that the coefficient for the interaction term between the monetary policy instrument

Table 4. Too Low for Too Long—Saturated Specification

Dependent Variable: Zombie	(1)	(2)	(3)	(4)	(5)
Shadow Rate $\times$ Long	-3.78*** (0.51)	-2.58*** (0.56)	-3.12*** (0.43)	-3.55*** (0.39)	-1.64*** (0.57)
Controls	Yes	Yes	Yes	Yes	Yes
Time	No	Yes	Yes	No	No
Firm	No	No	Yes	Yes	No
Sector	No	No	No	Yes	No
Sector $\times$ Time	No	No	No	No	Yes
Country	No	No	No	Yes	No
Observations	433,151	433,151	433,151	423,851	433,151
R^2	0.095	0.097	0.56	0.56	0.10
R^2_{adj}	0.095	0.097	0.49	0.49	0.10
Note: Coefficients expressed in percentage points. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.					

and our “Long” measure is negative and significant in all settings. This supports the claim that keeping interest rates too low for an extended period increases the probability of a firm turning into a zombie firm.


3. Zombification and Economic Growth

3.1 *Impact on Growth*

The previous micro and macro analyses suggest that extended periods of low interest rates can indeed facilitate zombification, but it needs to be clarified if and how the rise of zombie firms would reduce economic growth. On the one hand, at the firm level, there is a negative correlation between GDP growth and the probability of being a zombie firm. On the other hand, this relationship is not always significant on the macroeconomic side, as seen in Table 2. To shed light on this matter, we estimate a Bayesian panel VAR on the three regional units of our sample corresponding to monetary unions, namely the United States, Japan, and the euro zone. We identify the structural shocks using a Cholesky ordering. The variables of the system are the region-specific zombie share, the CPI inflation rate, the output gap, and the shadow rate in that order, based on the idea that monetary policy shocks do not have contemporaneous effects on the other variables of the system, which is consistent with the existing literature on structural VAR (SVAR) models, including Bernanke and Mihov (1998), Christiano, Eichenbaum, and Evans (1999), and Primiceri (2005). It is worth noting that the results are robust to different Cholesky ordering. Figure 5 shows the endogenous variables of the panel VAR system for each country of our sample. Given the time coverage of the firm-level data set, we estimate the model for the period 1996–2021.

As outlined by Canova and Ciccarelli (2013), there exist many alternatives to the estimation of panel VAR models in a Bayesian setting with different degrees of complexity. As we are interested in capturing the cross-sectional differences in the impact of the increase in the share of zombies on economic activity, we use the algorithm of Jarociński (2010). His Gibbs sampler employs hierarchical priors and treats the common mean and covariances of the VAR coefficients as random parameters. We estimate 50,000 Markov-chain Monte

Figure 5. Endogenous Variables of the Panel VAR




Note: Each column of the figure corresponds to a cross-sectional regional unit of the panel VAR, while each row shows the variables of the system. The zombie share is the fraction of zombie firms over the total number of firms for each regional unit; the inflation rate comes from the FRED (Federal Reserve Economic Data) database of the Federal Reserve Bank of St. Louis for the U.S., from the Bank of Japan’s website for Japan, and from Eurostat for the euro area; the output gap is obtained from the IMF database; and the monetary policy rates are the shadow rates as estimated by Krippner (2013).

Carlo (MCMC) simulations and discard the first 20,000 as burn-in. Figure 6 shows the impulse response functions (IRFs) following a one-standard-deviation shock increase in the region-specific zombie share on itself, CPI inflation rate, output gap, and the shadow rate. We can gauge different responses among the panel units.

In particular, the effect of the increased zombification on the output gap is negative in all the regions, peaking at negative 0.9 pp in Japan, followed by 0.75 pp in the euro area, although it is statistically insignificant in the case of the U.S. This might be evidence that the Japanese and European experiences in keeping interest rates very low for an extended period resulted in the zombification of their


Figure 6. Impulse Response Functions—Zombie Share Shock



Note: Each column of the figure corresponds to a cross-sectional regional unit of the panel VAR, while each row shows the variables of the system. The solid blue line is the posterior median response of each variable to a one-standard-deviation shock increase to the zombie share, with the corresponding 90 percent confidence bands in shaded blue.

weak firms, which was more detrimental to growth than in the U.S. Note that the shadow rates significantly decrease in response to the shock, which implies the monetary policy authorities responded to the increase in the share of zombie firms with expansionary policies. It is also worth noting that the lower shadow rates in response to the shock are more persistent in Japan, while they recover to the pre-shock level relatively quickly in the U.S., potentially explaining the differences in the adverse effects of zombification in the three regions. Figure 7 further supports the evidence that Japan and the euro zone suffered the most from the rise of zombie firms. Each chart corresponds to the region-specific forecast error variance decomposition of the output gap resulting from the panel VAR estimation. The output gap explains most of its variance throughout the response horizon period in all three regions. However, while the zombie share explains

Figure 7. Forecast Error Variance
Decomposition—Output Gap




Note: Each chart shows the region-specific share of the variance of the output gap, respectively, explained by shocks in the zombie share in green, the output gap itself in dark blue, the inflation rate in light blue, and the shadow rate in red.

respectively around 6 percent of the variability of the output gap in the U.S., this contribution accounts for up to 30 percent in Japan and the euro zone. As noted earlier on the higher persistency of the IRFs, the contribution of the zombie share is, on average, higher for the Japanese economy in the selected horizon, followed by the euro zone. This finding suggests that Japan, followed by the euro zone, experiences the sharpest economic slowdown explained by the zombification of its corporate sector.

3.2 *Spillover Effects*

The above adverse effects of zombification on economic growth might come from both direct effects of the presence of zombie firms with worse performance and indirect effects through impact on the more profitable and healthy firms. We employ descriptive statistics computed from our firm-level data set and compute zombie shares and profit margin growth for each sector. Figure 8 compares the full-sample averages of the two variables, from which we can notice some degree of negative correlation between sector performance and the zombie share. The sectors with lower profit margin growth, like textiles, wearing apparel, leather, and the primary sector, seem to have a relatively higher number of zombies. At the same time, gas,

Figure 8. Sectoral Zombie Shares and Economic Performance



Note: For each sector, the red and the blue bars correspond to the average zombie share and the average profit margin, respectively.
Sources: Orbis database, and the authors’ calculations.

water, and electricity, experiencing the lowest sectoral zombie share, also have the highest average profit margin growth.

In order to ascertain this hypothesis, we test whether high shares of zombie firms had a negative impact on the healthier firms of the same sector. We consider a similar specification as in Acharya et al. (2019) and estimate the following panel fixed-effects regression:

$$y_{i,j,h,t} = \beta_1 \cdot Nonzombie_{i,j,h,t} + \beta_2 \cdot Nonzombie_{i,j,h,t} \cdot Sector\ Frac\ Zombie_{j,h,t} + \gamma \cdot X_{i,j,h,t} + u_{i,j,h,t}, \quad (3)$$

where dependent variables are the employment growth and the profit margin of firms; $Nonzombie_{i,j,h,t}$ is a dummy variable equal to one if a firm i , in country j , sector h , and time t is not a zombie according to our specification; $Sector\ Frac\ Zombie_{j,h,t}$ measures the sector-specific zombie share; and $X_{i,j,h,t}$ is a set of firm-level control variables including the ICR, loan growth, and the ratio of EBIT over total assets. In this specification, the coefficient β_2 estimates the impact that an increase in the sector-specific zombie share has on

the respective dependent variable of non-zombie firms belonging to the same sector, hence, if zombie firms produce negative spillovers to healthy firms. Table 5 shows that an increase in the sector-specific zombie share exerts a sizable reduction on both the employment growth and the profit margin of non-zombie firms operating in the same sector. By saturating our fixed-effects specification, we obtain an average 12 pp decrease on the employment growth and a 12.2 pp decrease on the profit margin of healthy firms. Our results are consistent with Banerjee and Hofmann (2018), who employ a similar measure of zombie firms. What might be in place is a potential transmission channel that leads to a more general resource misallocation problem. If prolonged periods of low interest rates increase the probability that unprofitable firms keep on “floating,” the market will indirectly deprive the healthier firms of investment possibilities, in turn leading to lower employment and profits. In this sense, monetary policy at very low level of interest rates could be considered a form of financial repression.

4. Financial Stability and Banking Supervision

Some existing literature points to a weak banking sector as a key factor facilitating zombie firms’ emergence and survival. Caballero, Hoshi, and Kashyap (2008) document that zombie lending in Japan in the 1990s was mainly driven by undercapitalized banks “ever-greening” to loss-making companies to avoid additional provisioning, which would reduce these banks’ capital adequacy ratios, and Blattner, Farinha, and Rebelo (2023) show that weak European banks directed cheap credit to nonviable firms in the aftermath of the GFC. At the same time, supervisory powers seem to play an important role in taming zombie lending, as documented by Bonfim et al. (2023), who find that supervisors’ inspections reduce the likelihood of banks to refinance zombie firms after the inspection.

4.1 Supervisory Power

In the absence of specific bank-level data, we leverage our extensive firm-level data set in conjunction with the World Bank’s Bank Regulation and Supervisory Survey (BRSS) to assess whether country-level supervisory power affects the probability of a firm becoming

Table 5. Spillovers to Non-zombie Firms

Dependent Variable: Employment Growth	(1)	(2)	(3)	(4)	(5)
Sector Frac Zombie $\times$ Nonzombie	-0.18*** (0.03)	-0.15*** (0.03)	-0.11** (0.05)	-0.06* (0.03)	-0.10* (0.05)
Observations	312,045	312,045	303,447	303,447	303,447
R^2	0.03	0.06	0.16	0.16	0.17
R^2_{adj}	0.03	0.06	0.03	0.03	0.04
Dependent Variable: Profit Margin					
Sector Frac Zombie $\times$ Nonzombie	-0.12*** (0.004)	-0.15*** (0.006)	-0.11*** (0.007)	-0.12*** (0.005)	-0.11*** (0.007)
Observations	300,812	300,812	292,659	292,659	292,618
R^2	0.29	0.30	0.64	0.64	0.65
R^2_{adj}	0.29	0.30	0.59	0.59	0.60
Controls	Yes	Yes	Yes	Yes	Yes
Time	No	Yes	Yes	No	No
Firm	No	No	Yes	Yes	Yes
Time $\times$ Country	No	No	No	No	Yes
Sector	No	No	No	Yes	Yes
Country	No	No	No	Yes	No
Note: Coefficients expressed in percentage points. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.					

a zombie, and how this varies alongside a loosening monetary policy. Following Barth, Caprio, and Levine (2001), we construct an updated version of the Official Supervisory Power (OSP) index. This variable is determined by yes or no responses to 16 questions asked to the countries' central banks and prudential authorities and captures the extent to which supervisory authorities can intervene and take corrective actions. Our index, spanning the period 1999–2016, has annual frequency and compiles information from all surveys released to date.¹¹ For this analysis, we estimate the following regression specification:

$$\begin{aligned} \text{Zombie}_{i,j,t} = & \beta_1 \cdot \text{Shadow}_{j,t} + \beta_2 \cdot \text{OSP}_{j,t} \\ & + \beta_3 \cdot \text{Shadow}_{j,t} \cdot \text{OSP}_{j,t} + \gamma \cdot X_{i,j,t} + u_{i,j,t}, \end{aligned} \quad (4)$$


where $\text{OSP}_{j,t}$ is the logarithm of the supervisory power index in country j at time t , and $X_{i,j,t}$ includes the same set of control variables as in our baseline specification. The coefficient β_2 estimates the impact of an increase in country-level supervisory power on the probability of a firm becoming a zombie, while β_3 estimates how the interaction between supervisory power and monetary policy rates affects this probability. Table 6 presents the results of this analysis with a comprehensive set of fixed effects. The first row aligns with our previous analyses, indicating a negative relationship between interest rates and the probability of firms becoming zombies. The second row presents the estimated coefficients of our OSP series, which are negative and significant, highlighting a negative correlation between supervisory power and the likelihood of a firm becoming a zombie. This finding confirms that increased supervisory stringency, through more frequent inspections and enhanced control powers, can mitigate the “evergreening” of loans to unprofitable companies. The third and final row illustrates that higher supervisory powers can partially counteract the adverse effects of expansionary monetary policy. However, since the coefficient for the interaction term is positive but lower than the main coefficient for

¹¹Note that, since the surveys are not released annually, we build the index based on the release dates and assume that supervisory power remains constant until a new survey is released. This approach aligns with the fact that prudential regulations change with low frequency, allowing us to avoid distorting the data by interpolating between two adjacent values.

Table 6. Official Supervisory Power

Dependent Variable: Zombie	(1)	(2)	(3)	(4)	(5)
Shadow Rate	-0.66*** (0.21)	-0.55*** (0.24)	-0.19*** (0.025)	-0.24*** (0.023)	-0.08 (0.24)
OSP	-0.039*** (0.001)	-0.039*** (0.001)	-0.021*** (0.001)	-0.011* (0.003)	-0.038*** (0.001)
Shadow Rate × OSP	0.003*** (0.001)	0.003*** (0.001)	0.008*** (0.001)	0.008*** (0.001)	0.002*** (0.001)
Controls	Yes	Yes	Yes	Yes	Yes
Time	No	Yes	Yes	No	No
Firm	No	No	Yes	Yes	No
Sector	No	No	No	Yes	No
Sector × Time	No	No	No	No	Yes
Country	No	No	No	Yes	No
Observations	330,884	330,884	321,540	321,540	330,884
R^2	0.10	0.10	0.58	0.58	0.11
R^2_{adj}	0.10	0.10	0.52	0.52	0.10
Note: Coefficients expressed in percentage points. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.					

Figure 9. Monetary Policy (right-hand side) and Bank Lending Standards (left-hand side) in the U.S.



Sources: Orbis database, IMF World Economic Outlook database, and the authors' calculations.

the shadow rate, decreasing interest rates remain the primary driver of company zombification, even with higher supervisory power.

4.2 Lending Standards

Lastly, we try to gauge if banks had leniency toward weak companies by analyzing trends in net changes in the standards and terms on bank loans to SMEs in the U.S. using the Senior Loan Officer Opinion Survey on Bank Lending Practices (SLOOS). As can be seen from Figure 9, there is a strong correlation between the shadow rate (with a two-year delay) and the loan standards in the U.S. in the periods when the shadow rate was reduced to low levels. Although low interest rates and government incentives would normally be expected to support credit growth, the persistent easing of loan standards during 2004–06, when financial imbalances were building ahead of the GFC, and during 2010–18, when the economy was recovering from the GFC, can arguably be interpreted as a sign of greater leniency toward weaker firms.

However, Favara, Minoiu, and Perez-Orive (2021) find that U.S. banks do not seem to offer credit on favorable terms to zombie firms. On average, they charge zombie firms higher loan spreads, require

more collateral, and offer shorter-maturity loans than to non-zombie firms. Moreover, banks appear to assess zombie firms as having a higher probability of default and systematically rate their credit ratings lower than those of other firms.

Overall, related financial-sector risks in the 21 countries studied appear manageable in the short-to-medium term. Banks have significantly improved their lending and loan classification practices, capital and liquidity adequacy ratios, and stress-testing capacity in the aftermath of the GFC, especially in the context of Basel III reforms. Recent stress tests conducted by banks, regulatory authorities, and the International Monetary Fund (IMF), as part of their respective Financial Sector Assessment Programs (FSAP), indicate that the financial systems in these countries are capable of withstanding significant but plausible shocks. Moreover, these countries have significantly improved their financial-sector supervision and crisis management and resolution frameworks, notably with respect to early warning systems and pre-insolvency frameworks, enabling them to identify adverse trends and manage crises if they occur. Indeed, the financial systems performed well during the COVID-19 pandemic, a once-in-a-century shock. However, the prolonged period of low interest rates poses challenges for financial markets in the medium-to-long term, especially in Japan, as noted in IMF (2017). Profitability pressures are even more intense for smaller institutions outside urban regions in Japan due to declining populations in these areas. Furthermore, the prolonged low domestic interest rate environment is pushing life insurers to turn to foreign investments to provide the yield needed to meet interest guarantees on old policies. These developments imply risks that need to be continuously monitored.

5. Conclusions

While expansionary policies that boost economic activity and financial indicators of firms are necessary to smooth economic cycles, extended periods of low monetary policy rates may have undesirable adverse effects. Our empirical analyses confirm that recessions and slow growth are critical factors in the increase in the number of zombie firms. Accordingly, expansionary policies may be used in such cases if the central bank is credible and there are no inflationary

pressures. However, we also find that an extended low interest rate environment facilitates the survival of zombie firms, adversely affecting non-zombie firms through resource misallocation. This implies a trade-off between conducting countercyclical monetary policy and using expansionary tools for long periods, which may lead to a combination of low interest rates, low growth, and higher financial vulnerability. Such a trade-off is not a concern when most countries are tightening their monetary policy stance, but policymakers should be mindful of it during future recessions. Our findings suggest that Japan and, to a lesser extent, the euro area experienced a more significant economic slowdown partly due to the increase in the number of zombie firms. This raises concerns about the sustainability of keeping monetary policy rates low for extended periods, especially in countries with lackluster growth. We also find that stronger supervisory power is associated with a lower likelihood of firms becoming zombies. However, when combined with a low interest rate environment, this mitigation effect diminishes slightly. Despite the important reforms since the GFC, there remain risks to financial stability from extended periods of low monetary policy rates, especially if these need to be followed by policy tightening later, which must be closely monitored.

Appendix. Robustness Exercises

This appendix includes a set of robustness exercises to strengthen the paper's main results. First, we check the robustness of our main findings by departing from our preferred identification scheme of a zombie firm in our data set. Table A.1 shows the results of different specifications of Equation (3), where a firm is classified as a zombie if it is at least 10 years old and its ICR is lower than one for at least three consecutive periods, compared with the two periods in our baseline approach.

The results are similar both from a qualitative and quantitative perspective, and the economic intuitions drawn in Section 3 remain valid. The correlation between the shadow rate and the probability of becoming a zombie firm depends on the economic environment, with expansionary policies being associated with lower probabilities of a firm being a zombie under the ZLB and being associated with

Table A.2. Effect on the Probability of Being a Zombie Firm during Long Periods of Low Interest Rates (robustness exercise)

Dependent Variable: Zombie Share	Linear Probability Model		Probit Model (<i>mfx</i>)	
	(1)	(2)	(3)	(4)
Shadow Rate	−0.22*** (0.03)	−0.22*** (0.03)	−0.04*** (0.01)	−0.04*** (0.01)
Long	1.71*** (0.27)	1.43*** (0.27)	0.81*** (0.11)	0.68*** (0.11)
Shadow Rate × Long	−0.55* (0.30)	−0.28 (0.30)	−0.41** (0.19)	−0.29* (0.19)
Recession		0.55*** (0.08)		0.24*** (0.04)
SME	0.62** (0.26)	0.62*** (0.26)	0.20*** (0.07)	0.20*** (0.06)
GDP Growth	−5.69** (0.84)	−2.96** (0.95)	−2.68*** (0.41)	−1.40*** (0.46)
ROA	−0.29*** (0.01)	−0.29*** (0.01)	−0.21*** (0.01)	−0.21*** (0.01)
Profit Margin	−0.18*** (0.01)	−0.18*** (0.01)	−0.04*** (0.01)	−0.04*** (0.01)
Firm FE	Yes	Yes	No	No
Observations	433,151	433,151	433,151	433,151
Individual Firms	63,488	63,488	63,488	63,488
<i>R</i> ²	0.07	0.07		
Note: Coefficients expressed in percentage points. Standard errors in parentheses. * <i>p</i> < 0.10, ** <i>p</i> < 0.05, *** <i>p</i> < 0.01.				

a higher probability of being a zombie in the other periods. A difference is that the coefficient for the lagged shadow rate in column 4 here is very low and statistically insignificant. Second, in Table A.2, we replicate the estimation of Equation (2) with the zombie firm identification used in Table 5. As in the previous exercise, we can draw conclusions similar to our baseline estimation. Keeping the interest rates low for an extended period is correlated with a higher probability of a firm being a zombie.

References

- Acharya, V. V., M. Crosignani, T. Eisert, and S. Steffen. 2022. "Zombie Lending: Theoretical, International, and Historical Perspectives." *Annual Review of Financial Economics* 14: 21–38.
- Acharya, V. V., T. Eisert, C. Eufinger, and C. Hirsch. 2019. "Whatever It Takes: The Real Effects of Unconventional Monetary Policy." *Review of Financial Studies* 32 (9): 3366–3411.
- Adalet McGowan, M., D. Andrews, and V. Millot. 2018. "The Walking Dead? Zombie Firms and Productivity Performance in OECD Countries." *Economic Policy* 33 (96): 685–736.
- Banerjee, R., and B. Hofmann. 2018. "The Rise of Zombie Firms: Causes and Consequences." *BIS Quarterly Review* (September): 67–78.
- Barth, J. R., G. Caprio, and R. Levine. 2001. *The Regulation and Supervision of Banks Around the World: A New Database*, Vol. 2588. World Bank Publications.
- Bernanke, B. S., and I. Mihov. 1998. "Measuring Monetary Policy." *Quarterly Journal of Economics* 113 (3): 869–902.
- Black, F. 1995. "Interest Rates as Options." *Journal of Finance* 50 (5): 1371–76.
- Blattner, L., L. Farinha, and F. Rebelo. 2023. "When Losses Turn into Loans: The Cost of Weak Banks." *American Economic Review* 113 (6): 1600–1641.
- Bonfim, D., G. Cerqueiro, H. Degryse, and S. Ongena. 2023. "On-site Inspecting Zombie Lending." *Management Science* 69 (5): 2547–67.
- Caballero, R. J., T. Hoshi, and A. K. Kashyap. 2008. "Zombie Lending and Depressed Restructuring in Japan." *American Economic Review* 98 (5): 1943–77.
- Canova, F., and M. Ciccarelli. 2013. "Panel Vector Autoregressive Models: A Survey." In *VAR Models in Macroeconomics—New Developments and Applications: Essays in Honor of Christopher A. Sims*, ed. T. B. Fomby, A. Murphy, and L. Kilian. Emerald Group Publishing Limited.
- Christiano, L. J., M. Eichenbaum, and C. L. Evans. 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" In *Handbook of Macroeconomics*, Vol. 1, Part A, ed. J. B. Taylor and M. Woodford, 65–148. North-Holland.

- Corrado, L., S. Grassi, and E. Minnella. 2021. "The Transmission Mechanism of Quantitative Easing: A Markov-Switching FAVAR Approach." CEIS Research Paper No. 520.
- Diez, F. J., J. Fan, and C. Villegas-Sánchez. 2019. "Global Declining Competition." IMF Working Paper No. 19/82.
- Favara, G., C. Minoiu, and A. Perez-Orive. 2021. "US Zombie Firms: How Many and How Consequential?" Technical Report, Board of Governors of the Federal Reserve System.
- Giannetti, M., and A. Simonov. 2013. "On the Real Effects of Bank Bailouts: Micro Evidence from Japan." *American Economic Journal: Macroeconomics* 5 (1): 135–67.
- Harrison, R. 2012. "Asset Purchase Policy at the Effective Lower Bound for Interest Rates." Working Paper No. 444, Bank of England.
- . 2017. "Optimal Quantitative Easing." Working Paper No. 678, Bank of England.
- Hohberger, S., R. Priftis, and L. Vogel. 2019. "The Macroeconomic Effects of Quantitative Easing in the Euro Area: Evidence from an Estimated DSGE Model." *Journal of Economic Dynamics and Control* 108 (November): 1–25.
- Ikedá, D., S. Li, S. Mavroidis, and F. Zanetti. 2024. "Testing the Effectiveness of Unconventional Monetary Policy in Japan and the United States." *American Economic Journal: Macroeconomics* 16 (2): 250–86.
- International Monetary Fund. 2017. "Japan: Financial System Stability Assessment." IMF Country Report No. 17/244.
- Jarociński, M. 2010. "Responses to Monetary Policy Shocks in the East and the West of Europe: A Comparison." *Journal of Applied Econometrics* 25 (5): 833–68.
- Joyce, M., D. Miles, A. Scott, and D. Vayanos. 2012. "Quantitative Easing and Unconventional Monetary Policy—An Introduction." *Economic Journal* 122 (564): F271–F288.
- Krippner, L. 2013. "Measuring the Stance of Monetary Policy in Zero Lower Bound Environments." *Economics Letters* 118 (1): 135–38.
- Krishnamurthy, A., and A. Vissing-Jorgensen. 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." Technical Report, National Bureau of Economic Research.

- Lemke, W., and A. Vladu. 2017. "Below the Zero Lower Bound: A Shadow-Rate Term Structure Model for the Euro Area." ECB Working Paper No. 1991.
- Lombardi, M. J., and F. Zhu. 2018. "A Shadow Policy Rate to Calibrate US Monetary Policy at the Zero Lower Bound." *International Journal of Central Banking* 14 (5): 305–46.
- Mavroeidis, S. 2021. "Identification at the Zero Lower Bound." *Econometrica* 89 (6): 2855–85.
- Minsky, H. P. 1977. "The Financial Instability Hypothesis: An Interpretation of Keynes and an Alternative to 'Standard' Theory." *Challenge* 20 (1): 20–27.
- Mouabbi, S., and J.-G. Sahuc. 2019. "Evaluating the Macroeconomic Effects of the ECB's Unconventional Monetary Policies." *Journal of Money, Credit and Banking* 51 (4): 831–58.
- Peek, J., and E. S. Rosengren. 2005. "Unnatural Selection: Perverse Incentives and the Misallocation of Credit in Japan." *American Economic Review* 95 (4): 1144–66.
- Primiceri, G. E. 2005. "Time Varying Structural Vector Autoregressions and Monetary Policy." *Review of Economic Studies* 72 (3): 821–52.
- Turner, P., A. Rebucci, J. Zhou, and P. Monnin. 2019. "The Risks and Side Effects of UMP: An Assessment of IMF Views and Analysis." Background Paper No. 19-01/05, Independent Evaluation Office, International Monetary Fund.
- Vayanos, D., and J.-L. Vila. 2009. "A Preferred-Habitat Model of the Term Structure of Interest Rates." Technical Report, National Bureau of Economic Research.
- Wu, J. C., and F. D. Xia. 2016. "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound." *Journal of Money, Credit and Banking* 48 (2–3): 253–91.