

Is Monetary Policy Transmission Heterogeneous across Euro Area Countries and over Time? A Reassessment

Agnès Bénassy-Quéré¹, Matthieu Bussière², Thaïs Massei³
& Arthur Saint-Guilhem⁴

August 2026, WP #1062

ABSTRACT

We assess the degree of heterogeneity in monetary policy transmission (MPT) across euro area countries for the period 2000–2025. Using monthly local projections, we estimate the effects of monetary policy on a broad set of transmission variables at both euro area and national levels. First, we find limited heterogeneity of output and inflation responses to monetary policy shocks, despite asymmetric responses of mortgage market-related variables. Second, key structural differences, such as households' indebtedness, debt maturity, interest rate rigidity and sectoral composition, do shape heterogeneities in MPT across euro area countries, according to our results. Considered jointly, though, these differences partly offset one another, leading to a relatively homogeneous transmission of monetary policy. Third, a monthly FAVAR estimation confirms and broadens our local projections results: asymmetry remains contained for output and inflation, but it is higher for sovereign spreads, food prices, credit variables and our consumption proxy. Finally, a rolling-window estimation of the FAVAR model shows that time-varying heterogeneity in MPT is characterized by temporary and crisis-driven divergences that consistently revert to a low baseline, reflecting key monetary policy interventions rather than deeper structural economic divergences. Unconventional monetary policies play a central role in this pattern: they tend to reduce country divergences during crisis periods, impacting more strongly the countries most affected, hence operating, by design, as a heterogeneous policy shock according to our main metric.

Keywords: Monetary Policy Transmission, High-Frequency Identification, Local Projections, FAVAR

JEL classification: C32, C38, F45, E52, E31

*We would like to thank Jean Barthélemy, Christiane Baumeister, Tobias Broer, Edouard Challe, Alex Clymo, Catherine Doz, Guillaume Gaulier, Marek Jarociński, Stéphane Lhuissier, Barbara Rossi, Jamel Saadaoui for their comments and suggestions as well as all the participants of the internal Paris School of Economics Macro Workshop. We are particularly grateful to Georgios Georgiadis and Gregor von Schweinitz for their insightful discussions. Any views expressed herein are solely the responsibility of the authors and do not necessarily reflect those of the Banque de France and the Eurosystem.

¹ Banque de France and Paris School of Economics, Agnes.BENASSY@banque-france.fr

² Banque de France, Matthieu.BUSSIÈRE@banque-france.fr


³ Banque de France and Paris School of Economics | Paris 1, Thais.MASSEI@banque-france.fr

⁴ Banque de France and European Central Bank, Arthur.SAINT-GUILHEM@banque-france.fr

NON-TECHNICAL SUMMARY

When the European Central Bank raises or cuts interest rates, does every country in the euro area react similarly? The question is as old as the euro project itself: a monetary union is “optimal” only if cross-country divergences in the transmission of shocks are contained, so that its costs do not outweigh its benefits. The ECB itself has repeatedly acted on this concern, designing successive unconventional tools to prevent fragmentation within the euro area and ensuring effective monetary policy transmission (MPT) across euro area countries. This matters all the more because a few key structural features, such as wage-setting frameworks, lending markets characteristics or industry specialization, remain heterogeneous across countries and fiscal risk-sharing is still limited. The degree of heterogeneity in euro area MPT is therefore expected to be relatively high. Yet our results tell a different story.

Figure 1. Responses of output, prices and unemployment, one year after a monetary policy shock.



Notes: The blue dots show country-specific responses one year after a monetary policy shock, estimated using country-by-country local projections. The red dots show the euro area effect estimated on euro-area-wide data, while the green dots show the average response across the eleven countries obtained from a panel local projection. Vertical bars represent the 68% confidence intervals.

We reassess this question across a wide range of variables for the eleven founding euro area members between 2000 and 2025, using two complementary approaches: local projections and a Factor Augmented VAR model. Both approaches lead to convergent conclusions: heterogeneity in MPT is lower than previously assumed, especially for output and prices. Asymmetries do exist, but they lie elsewhere, in credit quantities, in our monthly consumption proxy, sovereign spreads and food prices.

Second, we study whether key structural characteristics contribute to observed heterogeneities in MPT. Household indebtedness, debt maturity, the prevalence of fixed versus adjustable-rate loans and sectoral composition each shape how national inflation and output respond to a common monetary shock: taken individually, each of them points toward greater divergence in MPT. But considered jointly, their effects tend to offset one another: high exposure to variable-rate mortgages, for instance, can be offset by comparatively low household indebtedness or a large share of long-term loans. This offsetting is unlikely to be coincidental: national economic and financial structures may have evolved together, each compensating for the vulnerabilities created by the other. Thus, assessed one characteristic at a time, the euro area looks structurally primed for divergent transmission but assessed jointly, these effects largely cancel out, explaining why the responses of output and prices end up more homogeneous than expected.

This full-sample result, however, does not tell us whether heterogeneity has stayed stable over 2000-2025, or whether it has instead evolved. Tracking our heterogeneity measure through a rolling-window FAVAR estimation, we find no structural drift: divergence rises temporarily around major monetary policy interventions, such as the 2015 launch of large-scale asset purchases (APP) or the 2022-2023 tightening cycle, before always reverting to a low baseline. Unconventional monetary policy tools play a central role in

this pattern: designed to smooth MPT within the monetary union, they do so precisely by acting, by construction, as a heterogeneous shock, impacting most strongly the countries that are most affected, for instance by compressing long-term sovereign yields disproportionately where they were initially most elevated.

La transmission de la politique monétaire est-elle hétérogène entre les pays de la zone euro et au cours du temps ? Une réévaluation

RÉSUMÉ

Nous évaluons le degré d'hétérogénéité de la transmission de la politique monétaire entre les pays de la zone euro sur la période 2000–2025. À partir de projections locales mensuelles, nous estimons les effets d'un choc de politique monétaire sur un large ensemble de variables de transmission, aux niveaux de la zone euro et national. Premièrement, nous constatons une hétérogénéité limitée des réponses de l'activité et de l'inflation aux chocs de politique monétaire alors que les variables associées au crédit aux ménages présentent des asymétries plus prononcées. Deuxièmement, des différences structurelles clés entre pays, telles que le niveau d'endettement des ménages, la maturité de la dette, la rigidité des taux d'intérêt et la composition sectorielle de l'économie, influencent effectivement les réponses nationales à un choc monétaire commun. Toutefois, lorsqu'elles sont considérées conjointement, ces caractéristiques se compensent, expliquant la transmission relativement homogène de la politique monétaire au niveau agrégé. Troisièmement, l'estimation mensuelle d'un modèle FAVAR confirme et élargit les résultats obtenus par projections locales : l'asymétrie reste modérée pour l'IPCH et le PIB réel, mais s'avère plus marquée pour les spreads souverains, les prix alimentaires, les variables de crédit ou l'indicateur de consommation. Enfin, une estimation du modèle FAVAR en fenêtre glissante montre que la dispersion des réponses nationales à un choc de politique monétaire varie au cours du temps et se caractérise principalement par des périodes de divergence temporaires liées aux épisodes de crise. Cette dispersion revient toutefois systématiquement à un niveau de référence faible, reflétant des interventions clés de politique monétaire plutôt que des divergences économiques structurelles profondes. Les politiques monétaires non conventionnelles jouent un rôle central dans cette dynamique : elles tendent à réduire les divergences entre pays en période de crise, en impactant plus fortement les pays les plus affectés, et opèrent ainsi, par construction, comme un choc de politique monétaire hétérogène, d'après notre principale métrique.

Mots-clés : transmission de la politique monétaire, identification à haute fréquence, projections locales, FAVAR

Les Documents de travail reflètent les idées personnelles de leurs auteurs et n'expriment pas nécessairement la position de la Banque de France. Ils sont disponibles sur publications.banque-france.fr

1 Introduction

Before monetary unification, in 1999, there was a heated debate on whether the euro area would be an *Optimum Currency Area* (OCA). The seminal paper of [Bayoumi and Eichengreen \(1992\)](#) suggested that, except for a small core of countries, both supply and demand shocks were far more asymmetric across European countries than across the regions of the United States, requiring distinct monetary policies. This diagnosis was however challenged along several lines. It was notably argued that, before monetary unification, the exchange-rate regime in Europe was not a flexible regime but a regime where national currencies were pegged to the Deutsche Mark, which means that even relatively large member states were importing German monetary policy. In a sense, monetary policy was a source of shocks for German partners, as much as a shock absorber ([Artis \(2003\)](#)).

Another area of discussion was whether the single monetary policy would be transmitted symmetrically across the members of the monetary union: structural heterogeneities could possibly turn symmetric shocks into asymmetric outcomes, due to asymmetric monetary policy transmission (MPT). The convergence process associated with euro area entry, such as the compliance with the Maastricht criteria, was intended to foster more harmonized macroeconomic adjustments and mitigate such risks. Although the launch of the euro initially encouraged greater alignment—particularly through the harmonization of the banking channel of MPT ([Angeloni and Ehrmann \(2003\)](#))—subsequent research has highlighted the persistence of asymmetries in country-specific responses to monetary policy shocks ([Barigozzi et al. \(2014\)](#); [Georgiadis \(2015\)](#); [Corsetti et al. \(2022\)](#)).

After a prolonged period during which monetary policy was either largely stalled (during the “Great Moderation”) or constrained by the effective lower bound, the 2022–2024 monetary policy cycle provides a unique opportunity to reassess asymmetries in MPT across the euro area. Indeed, between July 2022 and September 2023, the European Central Bank (ECB) raised policy rates by a cumulative 450 basis points in response to strong inflationary pressures. Then, starting in July 2024, policy rates were gradually lowered, in line with disinflation and the achievement of the inflation target.

The question of the degree of heterogeneity in MPT across euro area countries remains as relevant as it was at the inception of the euro for several reasons. First, there are many structural differences across euro area members, ranging from the characteristics of their lending markets (fixed vs flexible interest rates, loan maturity, borrower-based requirements, access to market funding, etc.) to industry specialization (manufacturing vs services) and wage setting frameworks (frequency, centralization, minimum wage, indexation rules). While differences are likely to create sources of heterogeneity in the transmission of monetary policy, these structural differences may also partially offset each other, with ultimately an ambiguous impact on MPT ([Bénassy-Quéré \(2025\)](#)).

In addition, while inflation had receded at the euro area level and in all countries in 2025, substantial cross-country differences in inflation rates persisted. For a central bank in charge of monetary policy for a currency area of 21 countries, it remains a policy relevant question to assess how the transmission of monetary policy might differ across countries and to what extent this could be an issue for effective monetary policy making. Moreover, over the past two decades, successive large shocks have affected member states unevenly, potentially generating divergences in monetary policy effects. Finally, the euro area is still characterized by limited fiscal risk sharing, departing from the canonical conditions of an OCA, which, in principle, may hinder its ability to stabilize heterogeneous economies. Whether these shortfalls have resulted in persistent differences in monetary policy responses over the history of the euro area remains an open question.

MPT is not expected to be perfectly symmetric across the members of the euro area. A single policy stance—which, by construction, targets aggregate euro-area conditions without explicitly minimizing cross-country dispersion—cannot fully accommodate country-specific stabilization needs. Some degree of asymmetry is therefore a fundamental and unavoidable welfare cost of monetary unification ([Corsetti \(2008\)](#)), and similar divergences are observed in federal countries such as the United States. Thus, together with the asymmetry of shocks, it needs to be balanced against the gains from the credibility of a single, independent monetary policy, on top of reduced uncertainty and transaction costs. This is the core of the OCA argument ([Mundell \(1961\)](#)). The viability of the union therefore depends on those divergences remaining contained. This is why, in this paper, we intend to reassess the asymmetry of MPT across the euro area.

We first measure the degree of heterogeneity in MPT by estimating both euro area and country-specific responses to a common monetary policy shock. Our analysis covers a broad set of transmission-related variables, including macroeconomic aggregates, borrowing costs and credit volumes, and focuses on the eleven founding countries of the euro area from January 2000 to December 2024. We estimate country-specific local projections à la [Jordà \(2005\)](#) using monetary policy shocks that we build from high-frequency data, following the approach of [Altavilla et al. \(2019\)](#) and [Jarociński and Karadi \(2020\)](#). This framework allows us to compare national estimates to several euro area benchmarks. We are also able to evaluate whether country-level responses are clustered around the aggregate or, alternatively, whether the euro area average masks substantial cross-country dispersion. While the local projection method is well established in the empirical literature, to our knowledge, no study has yet employed it to deliver a systematic and comparable assessment across both a broad set of euro area countries and variables. We implement a unified estimation strategy, applying a consistent specification of lags and controls across countries and variables, avoiding any ad hoc adjustments that could bias cross-country

comparisons. We then formally assess heterogeneity using the indicator proposed by [Corsetti et al. \(2022\)](#), which captures the average deviation of national estimates from the euro area mean.

For all the reasons previously mentioned, the degree of heterogeneity in euro area MPT is often expected to be high. Yet our findings point to a more nuanced picture. We find a low degree of heterogeneity in MPT for activity and inflation, although divergences tend to be larger for mortgage-related variables. Importantly, the relatively low dispersion in the responses of real GDP and inflation to a monetary policy shock can be interpreted as a conservative finding, given that local projection estimates typically display some instability across horizons and are particularly noisy at longer horizons. This analysis also offers new insights on which transmission channels exhibit greater uniformity across countries. The banking channel of monetary policy appears to operate in a broadly consistent manner across member states. However, quantity-based variables (credit volumes) display greater cross-country heterogeneity than price-based variables (lending rates), while variables related to non-financial corporations exhibit a lower degree of heterogeneity compared to those related to households. The relatively homogeneous transmission to borrowing costs suggests that the intertemporal substitution channel of monetary policy works in a broadly consistent manner across countries, contrary to the income channel of monetary policy, whose effectiveness is conditioned on certain country characteristics such as the level of debt maturity or the share of adjustable rate loans. Finally, this replicable and adaptable approach provides a practical and policy-relevant tool for tracking and monitoring cross-country MPT across the euro area.

Second, we study whether key structural characteristics contribute to observed heterogeneities in MPT. Using both a state-dependent panel local projection approach à la [Ramey and Zubairy \(2018\)](#), where dummies capture different regimes, and a specification interacting the monetary policy shock with (continuous) industry specialization and leverage structure, we find that a higher share of long-term loans and a lower share of variable rate loans both reduce and delay the impact of a monetary policy shock on inflation and GDP. To the contrary, a high level of households indebtedness and a high ratio of manufacturing over services in the value added both amplify the effect of monetary policy on GDP. However, when considered jointly, the impact of these country-specific characteristics on MPT appears more nuanced. Indeed, we show that the dampening effect on MPT of a reduced exposure to interest rate fluctuations observed in countries with a high share of long-term loans can be offset by a high share of flexible rate loans. A similar offsetting mechanism arises between these loan characteristics and the level of households indebtedness. Overall, these results indicate that, while individual structural features matter for an effective MPT, their combined influence at the macro level is less straightforward, underscoring the importance of also adopting a broader perspective that considers the joint effect of multiple characteristics. Doing so, this analysis helps reconcile the extensive evidence of non-linearities and state-dependency in MPT from the literature with the low dispersion of country-specific responses,

offering a reassuring conclusion: despite structural sources of asymmetries, MPT across euro area countries remains smooth.

Last, we examine if and how the degree of heterogeneity in MPT has evolved over time since the introduction of the euro, a question raised early by [Boivin et al. \(2009\)](#) but that has, to our knowledge, remained unanswered in the literature. We estimate a monthly Factor Augmented VAR (FAVAR) model in a fully Bayesian setting, following [Baumeister et al. \(2013\)](#). Our FAVAR incorporates a broad set of activity, price, financial as well as banking variables, for the euro area and its founding member countries. Estimating the FAVAR over the full 2000–2025 sample confirms that our local projection results on MPT heterogeneity are robust to a change of methodology: heterogeneity in MPT appears again particularly contained for activity and prices. Asymmetries do exist, but they lie elsewhere, in credit quantities, retail trade turnover, sovereign spreads and food prices. Then, to assess the evolution of cross-country heterogeneity over time, we re-estimate the same FAVAR specification over rolling windows and compute the cross-country standard deviation of the IRFs to a monetary policy shock to obtain a single heterogeneity measure per variable and time window. The choice of the FAVAR to conduct this analysis notably reflects the need for an estimation framework that remains stable across different sub-samples, a condition that local projections often fail to satisfy due to their notoriously high sensitivity to sample length and composition. Our results indicate that cross-country heterogeneity in MPT remains generally contained over time, but is punctuated by temporary crisis-induced surges. The latter mirror the dynamics of interest rates and inflation over the two decades and illustrate notably how unconventional monetary policy tools, designed to address country divergences and smooth transmission, operated precisely by affecting euro area countries unevenly in order to restore that smoothness.

The rest of the paper proceeds as follows. Section 2 proposes a literature review. After presenting the empirical framework, Section 3 provides both a visual and a quantitative assessments of the degree of heterogeneity in MPT across euro area countries. Section 4 examines the role of key structural and country-specific factors in shaping MPT. Section 5 analyzes the evolution of the degree of heterogeneity in MPT since the euro introduction. Section 6 concludes.

2 Literature review

Our paper contributes first to a large and well-established empirical literature examining heterogeneity in MPT across euro area countries. Early studies—often based on pre-euro or short post-euro samples—generally found that the introduction of the single currency fostered greater harmonization in MPT across countries. Employing a FAVAR framework, [Boivin et al. \(2009\)](#) show that while the compo-

nents of GDP continue to display non-uniform reactions to monetary policy across countries, aggregate output tends to respond more homogeneously. [Angeloni and Ehrmann \(2003\)](#) identify a structural break in banks' behavior and in the pricing of banking products around the euro's introduction, consistent with a harmonization of the banking transmission channel.

Although these studies document important changes in the transmission of monetary policy following the adoption of the euro, they also point to substantial remaining differences in how monetary policy shocks affect unemployment, prices and monetary aggregates across member states. Using single-country VAR models, [Ehrmann \(2000\)](#) finds stronger output responses in Germany than in Spain, France and Italy, while inflation reacts less in France and Italy than in Germany and Spain. [Angeloni and Ehrmann \(2003\)](#) note that the harmonization of non-bank transmission channels lagged behind, with early post-euro data showing little sign of convergence. [Boivin et al. \(2009\)](#) further demonstrates that the responses of long-term bond yields—particularly more pronounced in Italy and Spain than in Germany, France or the Netherlands—along with those of consumer prices, investment and consumption, still exhibit substantial cross-country heterogeneity. Among these, monetary aggregates appear least influenced by common euro area factors, reflecting structural national differences such as variations in household financial assets, wealth composition and the degree of interest rate pass-through.

More recent research confirms these findings. In the vein of [Barigozzi et al. \(2014\)](#) and using a dynamic factor model estimated between 2000 and 2016, [Corsetti et al. \(2022\)](#)—whose work is closest to ours—find persistent cross-country differences in the transmission of monetary policy to consumer prices, private consumption, unemployment and housing prices. These divergences appear concentrated around specific variables and are largely attributable to structural differences in mortgage (loan-to-value) and labor markets, two key sectors that have not fully converged across euro area economies. [Corsetti et al. \(2022\)](#) and [Barigozzi et al. \(2014\)](#) share the observation that remaining differences in the responses stem primarily from national characteristics and from a lack of integration of certain markets rather than from the ECB's conduct of monetary policy. Similarly, [Mandler et al. \(2022\)](#), using a multi-country BVAR for France, Germany, Spain and Italy, report that changes in policy rates have a stronger impact on real GDP in Germany and a weaker one in Spain, while price responses are strongest in Spain and weakest in Germany, reflecting structural features such as the relative weight of interest rate-sensitive sectors, labor market flexibility, export orientation and banking competition.

Against this background, a first contribution of our paper is to revisit the extent of heterogeneity in MPT using a longer sample that includes a particularly marked tightening cycle. In contrast with a large part of the existing literature, which tends to emphasize persistent cross-country asymmetries, we find that the overall level of heterogeneity is rather limited, especially for key macroeconomic aggregates such as output and inflation. Moreover, our results confirm the hierarchy across variables: heterogeneity

is lowest for GDP responses, followed by inflation while it is more pronounced for mortgage credit. As far as MPT is concerned, and conditionally on our sample, the role of structural differences across countries, while observable, may be more limited than commonly assumed.

Our study also connects to the literature that examines more specifically how monetary policy interacts with pre-existing specific structural and financial conditions—at the country, regional, sectoral, firm and household levels—thereby generating heterogeneous macroeconomic outcomes across and within countries. Among the country-specific factors, the structure of mortgage markets and more particularly the prevalence of adjustable-rate mortgages (ARM), in interaction with housing market features and households' saving behavior, have been shown to play a central role ([Calza et al. \(2013\)](#); [Wong \(2021\)](#)). Economies with a higher share of ARM tend to display stronger house price responses to monetary policy shocks ([Hedlund et al. \(2025\)](#)). Based on a local projections approach, [Pica \(2022\)](#) and [Sciacovelli \(2025\)](#) further demonstrate that the amplifying effect of ARM interacts with other characteristics: for the former, using the case of Spain, a high share of ARM combined with elevated homeownership rates increases households' direct exposure to interest rate fluctuations, while for the latter, the effect of ARM intensifies when a larger share of households is, as in Italy, liquidity-constrained, since fluctuations in mortgage payments more readily affect consumption in such context. Using Swedish micro data, [Flodén et al. \(2017\)](#) similarly document a cash-flow channel whereby monetary policy affects consumption more strongly for households that are more indebted and more exposed to ARM. [Luetticke \(2021\)](#), [Slacalek et al. \(2020\)](#) and [Cumming and Hubert \(2022\)](#) show respectively how heterogeneity in balance sheets (share of liquid versus illiquid assets), portfolio composition and exposure to aggregate fluctuations as well as households indebtedness condition the strength of the effect of monetary policy.

In the euro area, the structure of corporate financing also plays a key role in shaping MPT. Using a panel of French firms, [Alder et al. \(2023\)](#) document that conventional monetary policy shocks affect investment more strongly for firms relying on bank debt, whereas unconventional shocks have larger negative effects for firms more reliant on bond markets. Similarly, using firm-level data from ten euro area economies, [Holm-Hadulla and Thürwächter \(2021\)](#) finds that greater reliance on bond financing weakens the transmission of short-term policy rate shocks to real activity, as monetary tightening raises loan rates more than corporate bond yields, thereby limiting the overall increase in borrowing costs. According to [Georgiadis \(2015\)](#), countries where a large share of output stems from interest rate-sensitive sectors and with more real wage and/or fewer unemployment rigidities also tend to display a stronger transmission of monetary policy to real activity. Indeed, [Christoffel et al. \(2009\)](#) show that the level of nominal wage rigidity exerts a dominant influence on the propagation of monetary policy shocks, affecting both output and inflation dynamics. Sectoral composition further contributes to transmission asymmetries: for [Peersman and Smets \(2002\)](#), who document significant

cross-industry variation in the impact of monetary policy on GDP growth, these differences are linked to industry-specific features such as the durability of produced goods. Likewise, differences between capital-intensive manufacturing sectors and more labor-intensive services shape the strength of interest rate and cost-of-capital channels, with economies featuring a larger manufacturing sector, such as Germany, being more sensitive to both monetary tightening and easing ([Datsenko and Fleck \(2024\)](#) and [Hauptmeier et al. \(2025\)](#)).

Heterogeneity in transmission may also arise from different exposure to country-specific financial or external shocks. Using the global financial crisis, [Ciccarelli et al. \(2013\)](#) show that when countries, banks and firms are in weaker financial positions, they experience a more severe contraction in credit supply and sharper increases in financing costs, reinforcing asymmetries through both the bank lending—since weaker banks restrict credit supply more aggressively—and the borrower balance-sheet channel—as financially constrained or highly leveraged firms face greater difficulty accessing external finance. Thus, this financial fragility channel accounts for the heterogeneous and time-varying nature of MPT.

While our findings are consistent with this strand of the literature emphasizing the role of key structural characteristics in shaping MPT, we propose a novel interpretation that helps reconcile this evidence with the relatively low level of observed heterogeneity across countries. Specifically, we show that multiple country-specific characteristics, such as loan features, households indebtedness and sectoral composition, interact within countries in ways that can partially offset one another, thereby dampening cross-country differences in transmission.

Finally, our FAVAR-based analysis of the evolution of MPT heterogeneity builds on a well-established literature developing large-scale empirical frameworks for studying monetary policy effects. This line of research originates with [Bernanke et al. \(2005\)](#), who introduced the FAVAR approach to incorporate a broad set of macroeconomic and financial indicators into VAR analyses. In the euro area context, and alongside the dynamic factor model approach ([Corsetti et al. \(2022\)](#)), the FAVAR has proven particularly suitable ([Boivin et al. \(2009\)](#); [Ouerk et al. \(2020\)](#); [Barigozzi et al. \(2014\)](#)), as it allows to jointly estimate country-specific and area-wide IRFs while separating common euro area factors from national dynamics. This makes FAVAR models a natural tool for quantifying cross-country heterogeneity. For example, [Boivin et al. \(2009\)](#) employed a two-subsamples FAVAR estimation to assess the impact of the euro introduction on MPT. Subsequent extensions of the FAVAR literature ([Mumtaz \(2010\)](#); [Baumeister et al. \(2013\)](#); [Korobilis \(2012\)](#)) introduced time-varying frameworks, enabling the dynamic evolution of MPT to be studied in the United Kingdom and the United States. Thus, we contribute to this literature by explicitly studying the evolution of heterogeneity in MPT across euro area countries¹, a question largely

¹A natural extension of our framework would be to allow for time-varying parameters within the FAVAR, complementing the rolling window analysis presented in this paper.

unexplored. By documenting that heterogeneity has remained broadly contained over time and has primarily fluctuated with the monetary policy cycle, our study offers a new perspective on the stability of MPT within the euro area.

3 Assessing the level of heterogeneity across euro area countries

3.1 A local projection approach

3.1.1 High-frequency identification of monetary policy shocks and data

To identify exogenous changes in policy interest rates and construct our shock series, we rely on monetary policy surprises derived from high-frequency financial market data around ECB Governing Councils. All these high-frequency data are extracted from the *Tick History* LSEG Database. Specifically, we use minute-by-minute data from a large number of assets, including the Overnight Index Swap (OIS) rates, from January 2000 to December 2024, to replicate and extend the [Altavilla et al. \(2019\)](#) Euro Area Monetary Policy Surprise Database (EAMPD). Concretely, for each Governing Council day, we compute high-frequency surprises on the 1-month, 3-month, 6-month, and 1-year OIS maturities, as the difference in the median OIS rate over the 15 minutes before and the 15 minutes after the full monetary policy communication window. The full monetary policy communication window starts before the press release and ends after the press conference.²

The identification of monetary policy shocks using high-frequency surprises relies on two key assumptions. First, within the sufficiently narrow time window around a policy announcement, financial market movements are assumed to reflect only the central bank's communication rather than the market's feedback to other news or to contemporaneous economic developments. Second, monetary policy surprises are assumed to be unpredictable based on publicly available information prior to the announcement. That is, markets cannot systematically forecast the surprise using past macroeconomic data. This unpredictability implies that the surprises are exogenous, representing genuine shifts in monetary policy decisions rather than reactions to the current state of the economy. Consequently, these surprises, which capture unexpected changes in interest rates, can be regarded as valid instruments for identifying exogenous monetary policy shocks used in structural VAR and local projection analyses.³

Then, following [Jarociński and Karadi \(2020\)](#), we extract the first principal component from the four OIS surprises, obtaining a single indicator that captures the common monetary policy signal across maturities. The monetary policy surprises are then purged from the *central bank information effect*


²Except for the two exceptional dates (11/09/2001 and 08/10/2008), where only a statement was released.

³While some studies have questioned these identifying assumptions ([Bauer and Swanson \(2022\)](#)), addressing those concerns is beyond the scope of this paper.

through the sign restriction VAR approach developed by the authors. This isolates the pure monetary policy component from the surprise, i.e., the part reflecting genuine changes in the policy stance rather than shifts in the assessment of central bank's information about the economy. For example, when markets interpret a policy decision (such as a rate hike) as signaling that the central bank has received more optimistic news about future economic conditions, asset prices may move in the opposite direction to what would be expected from a pure policy tightening: a surprise policy tightening raises interest rates and reduces stock prices, while the positive central bank information shock raises both.

Finally, the monetary policy surprises are aggregated at a monthly frequency: in months without a Governing Council meeting, the shock is set to zero, while in months with multiple meetings, we sum the corresponding surprises. The robustness check, reported in Appendix A, confirms that our series of monetary policy surprises is fully comparable to the one commonly used in the empirical literature. The resulting series is standardized so that the shock is expressed in standard-deviation units. Thus, as shown in Figure 1, a one-standard-deviation surprise generates an immediate effect of 5.8 bps on the OIS1Y rate, with a peak response of 9 bps occurring three periods after the shock. For the OIS3M rate, it produces an impact effect of 5.6 bps and a peak response of 12 bps reached after four periods. This confirms that our shock is comparable to that in [Jarociński and Karadi \(2020\)](#), where a one-standard deviation shock leads to a 5 bps change in the one-year German government bond yield.

Figure 1: IRFs of the 3-month and 1-year OIS rates to a one-std monetary policy shock



Notes: Impulse responses of the 3-month and 1-year OIS rates to a one-standard-deviation monetary policy shock, estimated by local projections. The shock is identified using high-frequency monetary policy surprises and purged of the central bank information effect following [Jarociński and Karadi \(2020\)](#). The horizontal axis denotes the horizon in months; responses are expressed in percentage points. Shaded areas represent the one standard-deviation confidence intervals.

3.1.2 Econometric specifications

To evaluate the extent of cross-country heterogeneity in response to a monetary policy shock within the euro area, we rely on the local projection framework à la [Jordà \(2005\)](#). This framework notably offers

flexibility for also exploring nonlinearities in monetary policy transmission using state-dependent local projections (Section 4).

The main analysis is conducted at a monthly frequency, from January 2000 to December 2024, on the eleven founding countries of the euro area: Austria, Belgium, Finland, France, Germany, Ireland, Italy, Luxembourg, the Netherlands, Portugal and Spain. We restrict our analysis to this set of countries because local projections are highly sensitive to the number of available observations in the sample and estimation reliability increases with a longer sample. Unlike members that joined later, the founding countries are the only group for which we observe a long enough span of data as part of the euro area, with a sample comparable in length across countries. Including later entrants would produce estimates resting on different samples, spanning different economic events and yielding different degrees of precision. This would, in turn, alter comparability, as differences in sample length, rather than genuine differences in transmission, could drive the results. Second, our objective is to apply a single, unified specification (identical lag structure and controls) across all countries and variables, so as to ensure full comparability of the estimated responses. This guarantees that the cross-country heterogeneity we document reflects differences in country characteristics rather than differences in specification. We estimate the country-specific and EA responses to a one-standard deviation monetary policy shock (identified as described in section 3.1.1) of seven key variables used to assess monetary policy transmission: the *Harmonized Index of Consumer Prices* (HICP)⁴, real GDP, the unemployment rate as well as bank borrowing costs (of households for house purchase and of firms) and loan volumes for households and firms. Following [Holm-Hadulla and Pool \(2025\)](#), we exclude the COVID-19 period from our sample, as its inclusion generates abnormally large responses in real GDP and inflation—the only variables from our set whose dynamics appear significantly distorted by this episode.⁵

Data transformations. The dependent variable, denoted generically as y_{t+h} , takes different forms depending on the variable of interest. Following the recent recommendations of [Jordà and Taylor \(2025\)](#) and [Piger and Stockwell \(2023\)](#), we use a *long difference* specification to estimate the effects of a monetary policy shock on HICP and real GDP rather than estimating local projections in log levels. [Piger and Stockwell \(2023\)](#) highlights several advantages of estimating local projections in *long differences* when analyzing the effects of an exogenous shock that is identified (reduction of estimation bias, more accurate confidence intervals in small samples). In this case, the dependent variable y_{t+h} is transformed to log-difference, as an approximation of the corresponding growth rate:

⁴The HICP is the euro area’s standard consumer price index used by the ECB to define and assess price stability. Unlike national CPIs, which rely on different methodologies and are therefore not directly comparable across countries, the HICP is coordinated by Eurostat under a common methodology applied uniformly across all member states.

⁵Robustness checks comparing IRFs estimated with and without the COVID-19 period confirm that, once this episode is excluded, the magnitudes return to levels consistent with usual elasticities. This adjustment is justified, as the pandemic can be regarded as an exogenous shock—unlike other episodes over the sample period, such as the Global Financial Crisis, closely linked to underlying economic and financial dynamics.

$y_{t+h} = \frac{x_{t+h} - x_{t-1}}{x_{t-1}} \simeq \log(x_{t+h}) - \log(x_{t-1})$, where x_t denotes the variable (HICP or real GDP) in level. This choice implies that the estimated impulse responses represent the total cumulative percentage change in HICP and real GDP h months after a one-standard deviation monetary policy shock. For the unemployment rate, we apply the *long difference* directly to the level, so that the dependent variable is defined as $y_{t+h} = x_{t+h} - x_{t-1}$, where x_t denotes the unemployment rate in level. The estimated impulse responses therefore represent the total cumulative change in the unemployment rate, expressed in percentage points, h months after the shock. Loans to firms and households are expressed in year-on-year growth rate. Borrowing costs are kept in level. For each specification, the control variables are transformed accordingly to match the specification of the dependent variable. When necessary, series are deseasonalised.⁶ A table presenting each data series used and its transformation is presented in Appendix B.

Given the monthly frequency, the size of our sample and our object of interest —the effect of a monetary policy shock— we chose to include $p = 6$ lags⁷ for each variable and each specification. [Miranda-Agrippino and Ricco \(2021\)](#) go up to 12 lags at this frequency, while local projections run at a quarterly frequency usually include between 2 and 4.⁸ Our methodology is consistent across countries and variables. Thus, in each regression, irrespective of the dependent variable or of the country, the specification includes the six lagged values of the dependent variable, of the shock and of the control variables. Regarding the choice of the controls, we include country-specific and euro area controls. For real GDP, HICP and unemployment rate, we also add oil prices and the euro-dollar exchange rate as global controls. For the purpose of our analysis, we estimate the three complementary specifications of local projections, presented below:

- **Country-specific local projection:** this specification, which delivers our baseline results, gives us the country-by-country IRFs to a monetary policy shock, as well as the euro area estimates obtained by running the local projection on euro area-wide variables. For each country independently, we estimate:

$$y_{t+h} = \alpha^h + \beta^h \cdot \text{MPS}_t + \sum_{k=1}^p \phi_k^h \cdot X_{t-k} + \varepsilon_{t+h} \quad (3.1.1)$$

where y_{t+h} is the transmission variable of interest at horizon h , MPS_t denotes the monetary policy shock and X_{t-k} is the set of control variables (including p lags of the dependent variable and of the shock).

⁶Since country-specific HICP series are not provided deseasonalised, we seasonally adjust these data using the U.S. Census Bureau's X-13ARIMA-SEATS Seasonal Adjustment Program.

⁷Since our aim is to apply a single, common specification across both all variables and countries the number of lags was not selected on the basis of an information criterion, but rather guided by the literature and by the lag length typically used at the monthly frequency, relative to quarterly or annual estimations. Increasing the number of lags beyond six leaves our estimates broadly stable, so we retained this choice.

⁸Local projections have been shown to be robust to lag misspecification ([Jordà \(2005\)](#)).

- **Panel local projections (within estimation):** this specification depicts a panel local projection model estimated using a within (country fixed effect) estimator, controlling for time-invariant heterogeneity across countries. This gives us the average effect of a monetary policy shock based on the countries included in the sample. Depending on the countries included, we obtain two distinct euro area average effects from this specification: one based on 11 countries and another one based on 19 euro area countries.⁹

$$y_{t+h}^c = \alpha^{h,c} + \beta^{h,EA} \cdot MPS_t + \sum_{k=1}^p \phi_k^{h,EA} \cdot X_{t-k}^{EA} + \epsilon_{t+h}^c \quad (3.1.2)$$

- **Country-specific effects recovered from panel fixed effect estimation:** this panel local projection specification complements the two previous ones. By interacting the monetary policy shock with country dummies D_c , it allows to decompose the euro area average response β_h^{EA} obtained from specification (3.1.2), just above, into country-specific effects. In other words, this approach enables the estimation of both the average euro area response and individual country responses within a unified framework. In this setup, $\beta^{c_{ref}}$ captures the country-specific response of an arbitrary reference country (Austria, by alphabetical order), while for each other country c at horizon h , the response to a monetary policy shock is given by $\beta^{c,h} = \beta^{c_{ref},h} + \delta^{c,h}$. By construction, computing the unweighted mean of these country-specific effects yields the same value as β_h^{EA} estimated in the fixed-effects panel specification (3.1.2). Thus, this formulation provides a direct way to recover each country's responses from a single panel regression with fixed effects. National estimates obtained with this specification and with specification (3.1.1) are very close, though not identical because they do not capture the exact same dynamics (panel specification exploits cross-country correlations absent from the country-by-country set of regressions in specification (3.1.1)). Thus, the results of this specification, presented in Appendix D.1, are used for comparisons and robustness checks purposes.

$$y_{t+h}^c = \alpha + \sum_{c \neq c_{ref}} \gamma_c D_c + \beta^{h,c_{ref}} \cdot MPS_t + \sum_{c \neq c_{ref}} \delta^{h,c} (MPS_t \times D_c) + \sum_{k=1}^p \phi_k^{h,c} \cdot X_{t-k}^c + \epsilon_{t+h}^c \quad (3.1.3)$$

We compute heteroskedasticity robust standard errors following the recommendations found in the most recent studies by Jordà (2023),¹⁰ Inoue et al. (2024),¹¹ Montiel Olea et al. (2025) as well as Herbst and Johansson (2024) who all argue for the use of heteroskedasticity robust standard errors rather than heteroskedasticity and auto-correlation robust standard errors when conducting inference with

⁹All euro area countries are included, with the exception of Croatia and Bulgaria, as country data are incorporated only from the date of entry into the euro area (unbalanced setting). Given their recent entry, the time series for these two countries would be too short.

¹⁰"A simple way to obtain correct inference for the local projection is to add an additional lag as a regressor and then select a heteroskedasticity robust estimator to compute the standard errors—there is no longer a need to correct for serial correlation."

¹¹"Residual serial correlation in local projections can be dealt with and more recent developments show that it can be obtained rather easily using lag augmentation and heteroskedasticity robust methods."

local projections. For panel local projections, the baseline confidence intervals use heteroskedasticity-robust clustered standard errors in line with common practice in applied macroeconomics. As a robustness check, we also report the more conservative Driscoll–Kraay standard errors, following the recommendations of Inoue et al. (2024) for panel local projections.¹² The results are included in the appendix for comparison purpose. Finally, all the IRFs obtained from the local projection estimates are smoothed *a posteriori* using a moving average of order three following Falck et al. (2021). This post-estimation smoothing aims at reducing the horizon-to-horizon noise inherent to local projection estimates¹³ and so to improve the visual and interpretative clarity of the estimated responses without changing the dynamic pattern implied by the underlying estimated coefficients. The lower heterogeneity we document is not an artifact of the smoothing applied to the impulse responses. First, our results also hold with the unsmoothed estimates and are coherent across our two methods (LP and FAVAR). Smoothing is a late adjustment applied only for readability. Second, there is no mechanical effect of smoothing that could affect the measured heterogeneity: a three-horizon moving average applied to each country’s IRF separately neither changes the relative position of a country at a given horizon nor adds information, and can thus neither generate nor erase cross-country dispersion. Finally, all the statistical tests in the paper are computed from the raw estimates and standard errors, never from smoothed ones.

3.2 A visual inspection of the level of heterogeneity


In this section, we propose several ways of visually inspecting the degree of heterogeneity in monetary policy transmission. Before doing so, we first compare the different euro area responses of the variables of interest to a one-standard deviation monetary policy shock for four different specifications. These estimations differ by the underlying data chosen for the euro area and by the sample structure used: i) *the EA (aggregate data) case*, estimated from specification 3.1.1, displays responses based on area-wide data; ii) *the EA (panel - 11 countries) case*, estimated from specification 3.1.2, corresponds to the euro area average effect estimated with panel data of the eleven founding euro area members; iii) *the EA (panel - 19 countries) case*, also estimated from specification 3.1.2, corresponds to the euro area average effect estimated with panel data from all current member states, except Croatia and Bulgaria; iv) *The EA (GDP-weighted - all countries) case*, estimated from specification 3.1.3, reports a GDP-weighted average of

¹²These are robust to heteroskedasticity, autocorrelation and cross-sectional dependence in the error structure. It is this last property, robustness to cross-country correlation, that motivates this robustness exercise in our setting. Because the monetary policy shock is common to all countries, the residual of one country is likely correlated with those of the other countries at a given date. Clustering by country does not absorb this cross-sectional dependence, since it allows for correlation within a country over time but assumes independence across countries by construction. Moreover, cluster-robust inference relies on a large number of clusters and becomes less reliable with only a handful of countries. Driscoll–Kraay standard errors instead exploit the long time dimension of our panel: they remain valid as T increases, even with few cross-sectional units.

¹³Because each horizon h effect is estimated by a separate regression, the impulse response is a collection of independent point estimates rather than the output of a single recursive system, as in a VAR. With nothing tying adjacent horizons together, local projection responses are notoriously erratic.

the euro area response, constructed using data for all nineteen countries and weighting each national response by its share of euro area GDP.

Figure 2: Comparison of the EA responses to a 1-std monetary policy shock




Notes: The figure compares different euro area impulse responses to a one-standard-deviation monetary policy shock across four specifications. EA (aggregate data) uses area-wide series; EA (panel – 11 countries) relies on a balanced panel of the eleven founding members; EA (panel – 19 countries) uses a panel of all current member states except Croatia and Bulgaria; and EA (GDP-weighted – all countries) is the GDP-weighted average of the country-specific responses, each weighted by its share of euro area GDP. Shaded areas depict one-standard-deviation confidence bands based on clustered standard errors; the wider, more conservative bands based on Driscoll-Kraay standard errors are reported in Appendix C. No bands are reported for the EA (GDP-weighted – all countries), essentially shown for informational purposes, since the corresponding estimates are derived ex post by averaging the country-specific IRFs.

We observe that the resulting IRFs are strikingly close across all specifications, particularly for real GDP, unemployment and credit rates. For inflation, two groups emerge—one including the area-wide and 11-country panel estimates, and another comprising the 19-country panel estimation and the 19 countries GDP-weighted one. The green and purple lines lie close together, as expected, since the founding countries also tend to have the largest GDP shares in the euro area. Yet, even here, differences remain contained. Then, only the responses of credit volumes appear more dispersed across specifications. Overall, the closeness of the curves illustrates the stability and the robustness of the

euro area estimates across the different samples, suggesting that the later-joining countries do not materially alter the euro area response, so that the balanced panel of the eleven founding members remains representative of area-wide dynamics. Thus, it validates our decision to focus on the IRFs of the eleven original euro area members and to use these founding countries to further assess the level of heterogeneity in the euro area.

We continue our assessment of the level of heterogeneity in monetary policy transmission across euro area countries by comparing, in Figure 3, the IRFs across countries. It allows us to visually assess similarities in their magnitude (is the strength of the shock the same?), speed of adjustment (do some countries experience delays in transmission?), terminal effects after 18 months (are the outcomes, on a reasonable monetary-policy horizon, comparable?), and overall shape (do countries follow the same path after a shock?). Overall, while there are differences in magnitude or speed, they appear to be contained and the general pattern of responses looks comparable and coherent across countries, with few exceptions. We observe, that, at horizon 18, all country-specific IRFs of real GDP and HICP lie within the confidence interval of the euro area IRF, despite there is no statistical reason for this to occur, since all estimates are independently obtained with this specification. Differences in the pattern of transmission appear slightly more pronounced for unemployment rate responses since some countries exhibit a rapid and short-lived adjustment, while others display a more persistent effect over time.

Figure 3: Country-specific IRFs to a 1-std monetary policy shock



Notes: The figure plots the country-specific IRFs to a one-standard-deviation monetary policy shock, estimated with country-by country local projections, from specification (3.1.1). For comparison, the EA aggregate response (estimated on area-wide data) is shown together with its one-standard-deviation confidence band (blue shaded area); no bands are reported for the individual countries to preserve readability. For real GDP and HICP, the reported values represent the cumulative percentage change h months after the shock — for instance, a value of -0.15% at horizon 12 means that real GDP is 0.15% lower than it would have been in the absence of the shock. For the unemployment rate, the reported values represent the cumulative change in percentage points h months after the shock, relative to its pre-shock level — for instance, a value of 0.07 at horizon 12 means that the unemployment rate is 0.07 percentage points higher than it would have been in the absence of the shock.

In addition, Figure 4 allows us to focus more specifically on the dispersion of country-specific responses from the euro area estimates one year after the shock. By comparing these country-specific estimates with the euro area benchmarks, we can also gauge the degree of dispersion of national responses around the euro area ones and assess whether the average effect of ECB monetary policy is achieved at the cost of excessively tight conditions in some member states and overly loose ones in others. One

year after the shock, many country-responses are aligned with the euro area average for the three variables.¹⁴ In addition, the EA estimate lies within each country's confidence interval for almost all countries across the three variables. The HICP responses decline following a monetary tightening in every country, with no counterintuitive reaction. Very few IRFs display a price puzzle at the earliest horizons. One year after the shock, country estimates range from -0.22% in Belgium to -0.02% in Italy, averaging -0.13% with a cross-country standard deviation of 0.07%. This standard deviation hence roughly represents half the size of the average response, which indicates a moderate dispersion. The reactions of real GDP are again coherent across countries: output contracts in response to the shock in all cases. The degree of heterogeneity appears lower than for prices, in line with previous findings in the literature (Corsetti et al. (2022), Barigozzi et al. (2014), Boivin et al. (2009)). One year after the shock, country estimates range from -0.27% in Spain to -0.07% in Belgium, averaging -0.18% with a cross-country standard deviation of barely 0.07%. The latter represents around one third of the size of the average response, suggesting a moderate dispersion. For unemployment rate, one year after the shock, country estimates range from 0.02 percentage point in Germany to 0.13 percentage point in Spain, averaging to 0.08 percentage point, with a cross-country standard deviation of 0.04 percentage point, representing one half of the average response itself.


An important question is whether these country-specific IRFs primarily reflect underlying structural differences across member states. Some patterns are consistent with this interpretation. Germany's large GDP response aligns with its high manufacturing share in value added, a sector typically more interest-rate sensitive than services. Germany, Spain, Portugal and Italy which are characterized by infrequent wage adjustment and high wage stickiness¹⁵ show a stronger GDP drop than France and Belgium, where nominal wage rigidities are lower.¹⁶ While these examples suggest that structural features may help explain some cross-country differences, they also illustrate that no factor stands out to rationalize the full set of patterns observed. Indeed, Spain, where services account for a much larger share of value added, also exhibits a pronounced contraction of its activity, like Germany. It is important to acknowledge that these responses are likely the result of the combined influence of numerous factors playing in different directions and whose effects on monetary policy transmission, at the country level, may partially offset one another (see Section 4).

¹⁴The results of pairwise statistical tests of equality between country-specific IRFs are presented in Section 3.3.2, where we quantitatively assess the degree of heterogeneity.

¹⁵According to the OECD-AIAS Database, the average duration of wage contracts between 2000 and 2019 (the last year for which data are available) was 22,1 months in Germany, 22,8 months in Spain, 23,1 months in Portugal, and 30,6 months in Italy.

¹⁶According to the OECD-AIAS Database, the average duration of wage contracts between 2000 and 2019 was 12 months in France and 24 months in Belgium. If a 24-months wage contract duration is among the high ones, Belgium is however the only country among those under study whose collective agreements included automatic index or cost-of-living indexation clauses linking wages to prices.

Figure 4: Effects of a 1-std monetary policy shock, one year after the shock (h=12)




Notes: The figure reports the response of each variable to a 1-std monetary policy shock at horizon twelve for the eleven founding euro area members. Blue markers denote country-specific estimates; the red and green markers report the two euro area estimates: *EA aggregate data*, estimated from a local projection (specification 3.1.1) on area-wide series, and *EA panel - 11 countries*, obtained from a panel local projection (specification (3.1.2)). Vertical bars depict the one-standard-deviation confidence bands.

To better assess cross-country heterogeneity in monetary policy transmission beyond what can be inferred from individual IRFs, Figure 5 synthesizes country-specific responses with distributional statistics. For each horizon, we plot the median country response (black triangles), along with the 25th and 75th percentiles (blue dots), which describe the interquartile range (IQR) and therefore the dispersion of the central 50% of the distribution. The minimum and maximum responses (green triangles) capture the full extent of cross-country variation. Together, these elements provide a distributional overview of monetary policy transmission over time. The width of the IQR measures the degree of dispersion among the central mass of countries. It provides an assessment of the degree of heterogeneity that is robust to outliers. For all three variables, the IQR is particularly tight at short horizons and stays contained at longer ones, despite a widening. In addition, the distance between the extrema and the IQR side indicates whether any country deviates markedly from the others. For HICP and real GDP, even the most and least responsive countries do not exhibit opposite-signed or fundamentally different

dynamics after eight months. By contrast, the unemployment rate displays more extreme divergence, especially at longer horizons.

Figure 5: IRFs to a 1-std MP shock: min, max, median and interquartile range across countries



Notes: For each horizon, the figure summarizes the distribution of country-specific IRFs to a one-standard-deviation monetary policy shock across the founding euro area members. The width of the IQR captures the dispersion of the central mass of countries while the distance between the extrema and the bounds of the IQR signals the extent to which the most extreme countries deviate markedly from the rest.

We now focus on the banking variables that play a key role in monetary policy transmission in the euro area. Lending rates and credit quantities are typically driven by national structural characteristics that could contribute to make monetary policy transmission heterogeneous. As highlighted by [Chitu et al. \(2025\)](#), country-specific factors play a major role in shaping the bank-based external finance premium, as reflected in bank loan spreads, whereas they do not significantly affect the market-based external finance premium measured via corporate bond spreads.¹⁷ Our proxy of lending rates is the monthly interest rate on new loans granted by deposit-taking institutions (excluding the central bank) to non-financial corporations, on one side, and to households for housing purposes, on the other. As a measure of credit volumes, we consider the monthly annual growth rate of loans granted by monetary and financial institutions to non-financial corporations and to households and non-profit institutions serving households. Three key facts stand out from the four graphs presented in Figure 6.

First, the banking transmission channel of monetary policy appears to work in a rather harmonized way across the euro area countries presented here, as, for a given variable, IRFs share similar patterns and comparable magnitudes.

Second, prices of credit appear to exhibit a more homogeneous transmission pattern than quantities. Indeed, we see that borrowing costs, both for firms and households, tend to adjust more uniformly across countries than their corresponding annual growth rate of loans, which show greater dispersion,

¹⁷In the euro area, this stands in sharp contrast to bank loan spreads, which are strongly driven by country-level determinants, as recently documented by [Altavilla et al. \(2024\)](#). While banking remains a local activity, the euro-area corporate bond market is as integrated as that of the United States—contrary to conventional wisdom” ([Chitu et al. \(2025\)](#)).

more asymmetric patterns and several outliers in its trajectory.¹⁸ This suggests that the *intertemporal substitution channel* of monetary policy, operating through the price of credit, transmits smoothly across member states, despite different timing.

Third, variables related to non-financial corporations react more forcefully and rapidly, and along more homogeneous patterns, than those related to households. This observation stands out even more clearly when we turn to the IRFs derived from the panel estimation whose results are presented in Appendix D.1. Although the majority of mortgage credit rates follow a broadly similar pattern of response to a monetary policy shock, a few countries exhibit smoother adjustments in mortgage lending rates. France stands out as a notable exception: the slower increase in mortgage rates there typically reflects a regulatory mechanism (the usury rate) designed to protect households from excessive interest rate risk.¹⁹ Moreover, average rates on new mortgage loans tend to react less to monetary policy shocks in countries where long-term and fixed-rate mortgages dominate, such as France and Germany (See Figures in Appendix E.1). Such country-specific institutional features illustrate the well-known idiosyncrasies that continue to characterize mortgage markets across the euro area. In contrast, the credit rates charged to non-financial corporations are more clustered and exhibit a particularly similar hump-shaped response across countries, reflecting an even more uniform monetary transmission mechanism for firms than for households. This dichotomy between firm- and household-related variables is also reflected in the responses of credit quantity measures.


These differences largely come from the higher degree of integration in corporate credit markets compared to the persistent fragmentation of the household mortgage market in the euro area. Mortgage markets remain strongly shaped by national legal and institutional frameworks, including differences in mortgage contract rules such as the interest rate structures (fixed versus variable rates), debt maturities or early repayment rules. Moreover, cross-border borrowing remains particularly difficult for households, which contributes to less competitive borrowing conditions and, consequently, to greater heterogeneity in household borrowing cost dynamics across the euro area compared with firms. Unlike firms, households borrow almost exclusively from domestic banks, as foreign institutions face greater challenges in assessing individual creditworthiness, income stability and local property values. Finally, housing markets themselves are fundamentally local: there is no integrated EU-wide housing market. Property valuations, homeownership rates and borrowing practices differ markedly across member states, further reinforcing the fragmented nature of the mortgage credit

¹⁸The seemingly tight grouping of the IRFs for the growth rate of loans to households stems primarily from the Netherlands, whose pronounced outlier response expands the scale of the figure, making the IRFs of the other countries artificially close and homogeneous.

¹⁹Specifically, the legally mandated *usury rate*—a ceiling on the total cost of borrowing (including interest, insurance, term and risk premium) limits how quickly banks can raise rates. This cap is updated on the basis of a moving average, traditionally with quarterly revisions, although from February to December 2023 the frequency of revision was temporarily increased to monthly to allow for a more efficient transmission of policy rate changes.

market. In comparison, corporate credit markets are far more integrated. Corporate collateral is easier to value across jurisdictions since audited financial statements and credit ratings reduce informational asymmetries. Large and medium-sized firms can borrow from foreign banks or access international capital markets, which enhances cross-border competition and likely contributes to the coherence in loan price responses.

Figure 6: Impact of a 1-std monetary policy shock on banking transmission variables



Notes: The figure plots country-specific IRFs to a one-standard-deviation monetary policy shock, estimated with country-by-country local projections for the founding euro area members, together with the EA aggregate response and its one-standard-deviation confidence band. For credit rates, a value of 0.05 means the rate is 5 bps above baseline h months after the shock. For credit quantities variables, a value of -0.5 means that the year-on-year loan growth is 0.5 percentage point below baseline at horizon h . Unlike for macroeconomic variables, these values are not cumulated.

3.3 Quantitative measures of the degree of heterogeneity

3.3.1 Testing the equality of country-by-country estimates

We complement the visual-based evidence with two formal tests of equality of country-specific responses, focusing on the twelve-month horizon, a medium-term point at which the effect of monetary policy has largely materialized and that matters for its conduct. First, we perform pairwise tests of equality between the horizon-twelve coefficients of each pair of countries, $H_0 : \beta_i(12) = \beta_j(12)$, which yields 55 + 11 (with the EA estimate) unique pairs. Second, we report, for each variable, the share of pairwise tests (among the 55, excluding EA estimate here) for which the null of equality is rejected, providing a summary measure of dispersion. In this multiple testing framework, we base the reported

rejection shares on p -values adjusted for multiplicity using the Holm (1979) procedure, which controls the family-wise error rate (the probability of making even a single false rejection across the entire family of tests) under arbitrary dependence among the statistics.²⁰

Figure 7 summarizes the pairwise results in a heatmap reporting the 66 t -statistics, each given by:

$$t_{ij}(12) = \frac{\hat{\beta}_i(12) - \hat{\beta}_j(12)}{\sqrt{\widehat{\text{Var}}(\hat{\beta}_i(12)) + \widehat{\text{Var}}(\hat{\beta}_j(12)) - 2 \underbrace{\widehat{\text{Cov}}(\hat{\beta}_i(12), \hat{\beta}_j(12))}_{=0}}}$$

where the covariance term is omitted because the responses are estimated by independent²¹ country-by-country local projections, with countries on both axes and red dots indicating rejection of H_0 after the Holm correction. Reporting t -statistics lets the magnitude of each gap be read directly in standard-deviation units and the figure immediately reveals which clusters of countries have statistically indistinguishable responses and which stand apart. The test results are compatible with the contained degree of heterogeneity documented in other sections. For every variable, the overwhelming majority of pairwise statistics lie below the conventional threshold, indicating gaps between country responses that are small relative to their estimation uncertainty. Once the multiple-testing correction is applied, the number of rejections falls to zero for every variable.²²


Failing to reject equality is not equivalent to proving it: it may also reflect the limited power of the test, since local projection estimates are relatively noisy and the cross-section small. The question is then how large true heterogeneity would need to be for these tests to detect it at all. We therefore interpret the non-rejection as evidence that any residual differences are small relative to the precision with which they are measured.

²⁰ Conducting the 55 comparisons as a family inflates the probability of false rejections. A single test at the 5% level wrongly rejects a true null 5% of the time, so it behaves correctly with probability 0.95. If all country responses were truly equal and under independent test, the probability that all 55 behave correctly would be $0.95^{55} \approx 0.06$, so the probability of at least one spurious rejection would be $1 - (1 - 0.05)^{55} \approx 0.94$. A tolerable error rate for one test thus becomes near-certain false discovery across the family, which is why a raw count of rejections requires correction. Adjusting the p -values for multiplicity using the Holm (1979) procedure is the appropriate guarantee when the object of interest is a count of rejections.

²¹ The assumption of independence across country-specific estimates is, however, strong. We therefore repeat these tests using the responses from the panel framework (specification (3.1.3)), in which the covariance term is no longer zero. The results are presented in Appendix D.2. The conclusions are unchanged under both clustered and Driscoll-Kraay standard errors.

²² The Holm (1979) procedure works as follows. The p -values of the 55 pairwise tests are ranked from smallest to largest and each is compared to an increasingly lenient threshold: the smallest to $\alpha/55$, the next to $\alpha/54$, and the k -th to $\alpha/(55 - k + 1)$. One walks down the list and stops at the first p -value that exceeds its threshold, rejecting all hypotheses above it and retaining all those below. This controls the family-wise error rate—the probability of even a single false rejection across the whole family—under arbitrary dependence among the statistics. For inflation, the first pair faces the harshest bar, $\alpha/55 \approx 0.0009$, yet its p -value is already 0.15, so the procedure stops at the first comparison and rejects nothing.

Figure 7: Pairwise tests and joint test of equality of country-by-country responses at H=12



Notes: Each cell reports the t-statistic of the pairwise test of equality between the horizon-12 responses of two countries, $H_0 : \beta_i(12) = \beta_j(12)$. Responses are estimated with the country-by-country local projections from specification (3.1.1). Lighter cells correspond to smaller t-statistics; a value below 1.96 indicates that the two responses cannot be distinguished at the 5% significance level in an individual test. Red dots would flag pairs for which equality is rejected at the 5% family-wise level, after correcting for multiple testing via the Holm (1979) procedure; no such rejection occurs for any variable, which is why no red dot appears.

3.3.2 Measuring the dispersion of country estimates from the EA mean

To continue moving beyond visual inspection and quantifying cross-country heterogeneity in a more systematic manner, we compute a heterogeneity indicator, following the methodology proposed by Corsetti et al. (2022). In this context, heterogeneity now refers to the degree of dispersion of IRFs

around the average monetary policy response of the eleven euro area founding countries. For a given variable, this indicator measures the average distance of countries' responses from the euro area one. Specifically, for each horizon and variable, we calculate a coefficient of variation (CV), defined as the standard deviation of country-specific responses $\beta_{h,c}$, normalized by $|\beta_{h,EA}|$:

$$CV_h = \frac{\sqrt{\sum_c \frac{1}{N} * (\beta_{h,c} - \beta_{h,EA})^2}}{|\beta_{h,EA}|} \quad (3.3.1)$$

Normalizing by the size of the reference (the EA average) IRF makes this metric directly comparable across variables with responses of different scales. To assess the sensitivity of our estimated CV to each country, we implement a jackknife leave-one-out procedure. For each variable and horizon, we successively exclude one country at a time from the cross-sectional distribution of country-specific impulse responses and recompute the CV from the remaining ten countries. This yields eleven leave-one-out CV estimates per horizon, from which we derive a jackknife standard error following the standard formula. This procedure answers the question of whether the measured heterogeneity is robust to sample composition.²³ Our results are presented in Table 1. In our first set (so called *country-by-country LP*), the CVs are computed using the country-by-country local projection estimates already shown previously in Figures 3, 4, 5, 6 and obtained with specification 3.1.1 (from subsection 3.1.2). In this case, the euro area reference response $\beta_{h,EA}$ used for normalization is the unweighted average of the eleven country-specific responses estimated independently. In the second set of results (so called *country-by-country LP*), the CVs are based on the $\beta_{h,c}$ recovered directly from the fixed-effect panel local projections described in Specification 3.1.3 (from subsection 3.1.2). All the underlying IRFs corresponding to this specification are presented in Appendix D.1. In this second case, the euro area reference response $\beta_{h,EA}$ used for normalization is also the unweighted average of these eleven country-specific responses, estimated, here, jointly.

For each specification, we report the average CV over horizons 5 to 15,²⁴ a synthetic measure of heterogeneity that is less prone to being driven by explosive early CV outlier values, due to a near-zero EA mean mechanically inflating the ratio or due to an explosive standard deviation due to short-term price puzzles in some countries.²⁵

²³The jackknife confidence intervals reported here capture the sensitivity of the CV to sample composition, i.e. how the CV changes when individual countries are excluded from the sample. They do not, however, account for estimation uncertainty arising from the fact that each country-level impulse response is itself estimated with error. Addressing this alternative source of uncertainty would require, in our local projection setting, a wild bootstrap procedure applied at the regression stage, combined with bias-corrected (BC) confidence intervals rather than standard percentile intervals, given that the CV is a ratio statistic that is upward biased in finite samples. In practice, however, this correction is not operational here: while each point estimate is divided by a single denominator safely away from zero, the resampled cross-country means occasionally approach zero, producing explosive bootstrap draws. The bias correction relies on the fraction of bootstrap draws falling below the point estimate, but the upward bias pushes the estimate below the entire bootstrap distribution, driving that fraction to zero. The bias correction term becomes undefined and so is the confidence interval —at most horizons and particularly for the horizon-averaged CV of interest.

²⁴Similar ranking is obtained by averaging over other horizons.

²⁵In our local projection framework we chose to restrict the computation of our average CV to these medium horizons because, at shorter horizons, the effects of the monetary policy shock can be particularly small. Computing the CV increases, in this

Table 1: CV and bootstrapped confidence intervals by variable and horizon

Variable	Country-by-country LP			Panel LP		
	Average CV	LB*	UB*	Average CV	LB*	UB*
HICP	0.58	0.33	0.83	0.75	0.42	1.07
Real GDP	0.53	0.31	0.76	0.58	0.38	0.78
Unemployment rate	0.85	0.39	1.31	0.53	0.37	0.71
HH lending rate	1.03	0.46	1.59	0.52	0.29	0.76
Firms lending rate	0.48	0.27	0.69	0.30	0.21	0.39
Loans to HH	3.50	0.00**	14.58***	4.23	0.00**	29.25***
Loans to firms	0.86	0.61	1.11	0.51	0.31	0.71

* The jackknife standard error is computed as: $\hat{se} = \sqrt{\frac{N-1}{N} \sum_{c=1}^N (\widehat{CV}_{-c} - \overline{CV_-})^2}$ and confidence intervals are formed as $\widehat{CV} \pm t_{N-1, \alpha/2} \times \hat{se}$, where $t_{N-1, \alpha/2}$ denotes the $\alpha/2$ critical value of the t distribution with $N-1=10$ degrees of freedom. We report 90% confidence intervals. This interval captures only the sensitivity of the CV to which countries are included in the sample but is not informative of the precision of the country-specific estimates.

** The lower bound of the jackknife confidence interval is censored at 0 by construction, as the CV is non-negative. For highly heterogeneous variables, the leave-one-out CV exhibits substantial dispersion across drop iterations, yielding a jackknife standard error large enough that the raw lower bound turns negative and the floor binds. This should not be interpreted as potential evidence that $N=11$ is too small to provide sufficient precision to bound the magnitude of the CV from below. This pattern is not driven by any single country or horizon, suggesting this 0 lower bound does not reflect the influence of an outlier.

*** The large upper bound reflects the strong right-skewness of the CV distribution and sensitivity of this statistic to near-zero denominators.

The low degree of heterogeneity in the responses of real GDP and inflation, previously observed through visual inspection, is corroborated for both variables. The dispersion of country responses around the EA benchmark ranges from 53% to 58% of the EA reference for real GDP, and from 58% to 75% for HICP, in both cases well below the value of 1, at which dispersion would be 100% of the average EA response. While the homogeneity of real GDP responses has already been documented in the literature, our findings based on an extended sample which encompasses the sharp 2022-2024 monetary policy cycle, provide more novel evidence for inflation. There is no standard threshold of CV value that separates high from low heterogeneity: a CV of one merely marks the point at which dispersion equals the size of the average response, not a meaningful cutoff between high versus low CV. Therefore, a degree of homogeneity of a variable measured with this metric has to be also assessed relative to that of other variables, particularly those known to be more country-dependent by construction. Thus, in the baseline specification, the responses of the unemployment rate exhibit a slightly higher degree of heterogeneity than HICP and real GDP, as reflected in both a higher coefficient of variation and upper band. However, since this greater dispersion is not robust to the panel specification, the CV attached to the unemployment rate is best read as being of a similar order of magnitude as that of the other macroeconomic aggregates. The CV is particularly high for the growth rate of loans to

context, the likelihood of a near-zero euro area mean in the denominator, which can mechanically generate explosive values of the CV even when cross-country dispersion in the numerator remains modest. This may lead to an artificially high measure of heterogeneity driven by the metrics rather than by underlying economic differences. Excluding these horizons therefore helps avoid confounding these effects.

households. These findings align with the results of [Corsetti et al. \(2022\)](#), who, using a dynamic factor model approach to estimate the euro area and country-specific estimates, document deeper divergence in mortgage- and labor-market-related variables. Their analysis attributes this pattern to the slow convergence and low degree of integration of these markets since the euro introduction.

Moreover, the computation of the CVs confirms that the dispersion in loan volumes is substantially larger than that in borrowing costs, suggesting that cross-country differences in credit dynamics extend beyond differences in the price of credit. This divergence indicates that country-specific factors may play an important role in shaping the behavior of loan demand (national housing market characteristics) and loan supply (national regulations and macroprudential policies, loan-to-value limits, banking sector structures, competitiveness, bank balance-sheet constraints), on top and above the price of credit.

Finally, for both lending rates and loan quantities, heterogeneity is higher on the household side than on the firm side. This pattern confirms that monetary policy transmission through the household credit channel is significantly more heterogeneous than through the firm credit channel, reflecting, on the households side, heterogeneity in mortgage market structures and in housing markets, and, on the firm side, the larger firms' access to capital markets and cross-border financing which tends to harmonize credit conditions across countries.

Taken together, the evidence so far points to a contained degree of cross-country heterogeneity in the transmission of monetary policy to macroeconomic aggregates. This does not mean, however, that country-specific structural characteristics play no role. A low dispersion of the aggregate responses is compatible with monetary policy propagating quite differently across countries through individual transmission channels. What matters is whether these channel-level differences ultimately translate into sizeable aggregate dispersion, or whether they wash out, either because characteristics offset one another, or because the realistic spread of a given characteristic across euro area countries is too narrow to generate large divergences in the responses of aggregates. Structural features can thus each shape transmission at the country level while leaving the aggregate response relatively homogeneous. This is why we now turn to testing how country-specific differences shape the national monetary policy outcomes previously observed, opening the black box behind the aggregate impulse responses.

4 Explaining the potential sources of heterogeneities

4.1 Methodology

The literature on the sources of heterogeneity in monetary policy transmission is not new. Drawing on this extensive empirical and theoretical literature, we examine four structural features susceptible to shape cross-country differences in monetary policy transmission to real GDP and HICP:

- the maturity structure of households (HH) and non-financial corporations (NFCs) debt, proxied by the share of long-term loans (above 5 years) in the total amount of loans granted by monetary and financial institutions (MFI) excluding Eurosystem to HH and NFCs;²⁶
- the degree of interest rate sensitivity, proxied by the share of variable-rate loans in the total loans to NFCs and HH;²⁷
- the degree of households' indebtedness, proxied by the loans-to-income ratio;²⁸
- the sectoral composition, proxied by the ratio of manufacturing-to-services value added.²⁹

Country-by-country time series for these characteristics are presented in Appendix E.1. Identifying their individual effect helps to unpack the black box behind the IRFs previously presented and understand the forces driving transmission outcomes. Indeed, the fact that the IRFs of key macroeconomic aggregates display limited heterogeneity does not imply that no country-specific characteristics amplify or attenuate monetary policy transmission and identifying such characteristics remains a policy-relevant exercise.

To assess the role of these structural features, we exploit the flexibility offered by the local projection approach by interacting the monetary policy shock with structural characteristics, employing two complementary econometric specifications. The first interacts the monetary policy shock with a dummy variable that distinguishes low from high states for the different characteristics of interest, in a non-linear state-dependent local projection setting. The second interacts it with the structural characteristics measured as a continuous variable. Each specification carries distinct advantages and limitations for our question of interest and their combination provides a more robust basis for analysis than either approach taken in isolation. For this exercise, our sample goes again from January 2000 to December

²⁶The data used to compute the share are the national stocks of loans by maturity vis-a-vis euro area NFCs and HH (excluding general government) reported by MFI excluding ESCB from the ECB Balance Sheet Items Dataset.

²⁷The data used are the monthly share of variable rate loans —defined as floating rate loans or loans with an initial fixation period of up to one year— in the total loans to HH and NFCs from the ECB Risk Assessment Indicators Dataset.

²⁸The data used are the quarterly stock positions of loan liabilities for national HH and non-profit institutions serving HH, owed to the total economy expressed relative to annual gross disposable income, from the ECB Quarterly Sector Account Dataset.

²⁹The data used are the quarterly, seasonally and calendar-adjusted sector-by-sector gross value-added series, expressed as a percentage of GDP, sourced from Eurostat's national accounts (ESA 2010), NACE Rev. 2 classification.

2024 and includes, as long as data availability allows, all current euro area member states except Croatia and Bulgaria. By increasing the cross-sectional variation in the characteristics of interest, this broader coverage allows for a more precise estimation of the heterogeneous effects of interest and makes the allocation of countries across percentiles more informative. For each estimation, the lags and control variables are defined in the same way as the local projections presented in the preceding sections. The same variable transformations (*long-difference* applied to log of real GDP and HICP) as before are also applied. As in the previous section, the IRFs displayed are smoothed with a MA(3).

Specification 1. *Does the transmission of monetary policy vary across high versus low states of the characteristic?*

We follow the state-dependent fixed effect panel local projection framework à la [Ramey and Zubairy \(2018\)](#) and à la [Gonçalves et al. \(2024\)](#). The effect of a monetary policy shock is allowed to differ depending on whether the state of the characteristic in question is *high* or *low*. The dummy variable is constructed based on the median value of the corresponding structural characteristic across the whole sample. Thus, an observation is classified as belonging to the high state when its value exceeds the sample median (pooled median over countries and time) and to the low state when it lies below it.³⁰ This pooled median split ensures a balanced distribution of observations between the two regimes. Hence, this specification is the following:³¹

$$y_{t+h} = \alpha_h + \beta_h^a \cdot \mathbf{1}(D_{t-1} = 0) \cdot MPS_t + \beta_h^b \cdot \mathbf{1}(D_{t-1} = 1) \cdot MPS_t + \sum_{k=1}^p \phi_k^h \cdot X_{t-k} + v_{t+h} \quad (4.1.2)$$

The monetary policy shock is interacted with a lagged dummy D_{t-1} . In our setting, concerns about endogeneity arising from contemporaneous feedback are generally limited because the dummies are constructed from characteristics that are highly persistent over time. The main exception is for the level of household indebtedness, as a monetary policy shock may affect disposable income, which enters the denominator of our proxy for this characteristic. For each estimation, we perform, as a baseline, a Wald test of equality between the two state-dependent IRFs, based on clustered standard errors. Here again, we also report in [Appendix E.2](#) the confidence intervals and test results obtained under the more conservative Driscoll–Kraay standard errors, which account for serial correlation, heteroskedasticity and cross-sectional dependence across countries. Despite the substantially wider confidence bands they imply, most of our estimated economic effects remain statistically significant under this correction, pointing to the robustness of our results.

³⁰ As we assume that the marginal effects of the control variables on the dependent variable are independent from the dummy value, we choose not to interact the controls and the dummy.

³¹ Specification (4.1.2) is strictly equivalent to estimating

$$y_{t+h} = \alpha_h + \beta_h^a MPS_t + \beta_h^b D_{t-1} * MPS_t + \beta_h^c D_{t-1} + \sum_{k=1}^p \phi_k^h \cdot X_{t-k} + v_{t+h} \quad (4.1.1)$$

However, our chosen specification 4.1.2 allows for the direct estimation of separate coefficients for each state (β_h^a and β_h^b), along with their respective standard errors and confidence intervals. This, in turn, enables direct statistical inference and allows us to conduct a Wald test to assess the equality of the two states IRFs, at each horizon, in a straightforward way.

When the corresponding 68% level confidence intervals do not overlap, the difference between the two responses is visually evident; in such cases, the Wald test does not provide additional insight. However, when the intervals partially overlap, the Wald test becomes informative, as overlapping confidence bands do not necessarily imply the equality of the responses. A lower p-value from the Wald test indicates stronger evidence against the null hypothesis of equality.³²

This specification has several advantages. By interacting the monetary policy shock with a binary indicator, we assess whether a given characteristic influences the transmission of monetary policy in the euro area context without imposing any parametric restriction on the form of this effect. This contrasts with standard interaction models based on continuous characteristics, which implicitly assume a linear effect. The outcome is also straightforward to interpret: the specification yields two distinct impulse responses, one for the low state and one for the high state, whose difference can be directly tested via a statistical test. Furthermore, this approach is well suited to the nature of the characteristics under study, which are structural and therefore relatively stable over time, at least for long time ranges.

That said, this specification has limitations that motivate the complementary analysis conducted with Specification 2. First, for certain characteristics, individual countries may switch between high/low states over the sample period. This can either be due to a genuine structural break, as is the case for Spain which transitioned from a predominantly variable-rate regime to a fixed-rate regime around 2015, or simply because their values over the sample hover close to the median. We do not want the latter, these high-frequency transitions, to determine the identification of the effect of the characteristic. Second, several characteristics display a time trend over the sample period, and so does the resulting dummy variable. When low values are systematically concentrated at the beginning of the sample and high values at the end, or vice versa, the estimated differential in IRFs may reflect changes in the macroeconomic environment over time rather than the effect of the structural feature *per se*, confounding state-dependent differences in transmission with unwarranted time variation. Controlling for this trend is, however, not straightforward. Adding time fixed effects is precluded by the fact that the monetary policy shock is common to all countries in the panel. The use of a time-varying median instead of a pooled median split would remove the mechanical link between the dummy and a time trend, but it would change the economic question being asked since the dummy would no longer identify countries with structurally high or low characteristics, but merely those above the cross-sectional median at a given point in time (even if they are all low or all high), a conceptually distinct criterion. This distinction is further complicated by the fact that country rankings relative to one another are not fully stable over time, meaning that a given country may shift between the high and low state not because of any

³² Minor and isolated discrepancies may arise at certain horizons between the plotted confidence intervals and the Wald test outcomes. This reflects the fact that the IRFs and their associated confidence intervals displayed in the charts are smoothed whereas the Wald tests are conducted using the raw IRFs. Consequently, figures may occasionally suggest significance at the 68% confidence level, while the corresponding p-value slightly exceeds 0.32 and vice versa.

meaningful change in its own structural position, but simply because of shifts in the distribution of the characteristic across the panel.

Specification 2. *Does a higher level of the country-specific characteristic amplify or dampen monetary policy transmission?* We alternatively interact the monetary policy shock with the structural characteristic expressed as a continuous variable. The characteristic of interest $\tilde{Z}_{t-1}$ is the standardized version of Z_{t-1} , our characteristic of interest, transformed prior to estimation so that β_h^2 captures the effect of the monetary policy shock at the sample mean of the characteristic, while β_h^1 measures the change in this effect associated with a one-standard deviation increase in the characteristic.

$$y_{t+h} = \alpha_h + \beta_1^h \cdot MPS_t \cdot \tilde{Z}_{t-1} + \beta_2^h \cdot MPS_t + \beta_3^h \cdot \tilde{Z}_{t-1} + \sum_{k=1}^p \phi_k^h \cdot X_{t-k} + v_{t+h} \quad (4.1.3)$$

To facilitate an economically meaningful interpretation, the total effect of the monetary policy shock is evaluated at the 25th and 75th percentiles of the cross-country distribution of the characteristic. To obtain these percentiles, we first average each country's characteristic over the sample period, giving one value per country, and then take the 25th and 75th percentiles across these country averages. The impulse responses conditional on low and high levels of the characteristic are then given by $\hat{\beta}_h^1 + \hat{\beta}_h^2 \cdot z_{p_{25}}$ and $\hat{\beta}_h^1 + \hat{\beta}_h^2 \cdot z_{p_{75}}$, respectively.

This specification offers several advantages. First, this specification avoids an information loss inherent to specification 1: a binary split treats an observation just above the median the same as one far above, whereas the continuous interaction lets the magnitude of each country's deviation from the mean shape its contribution to the estimated effect. Second, by anchoring the evaluation points to the empirical distribution of structural characteristics across countries, we move beyond a purely qualitative assessment of whether a given characteristic shapes monetary policy outcomes or not and we provide a quantitative evaluation of its effect at values that reflect observed euro area cross-country heterogeneity. The resulting IRFs illustrate how monetary policy transmission operates in practice at representative points of the distribution, bringing economic plausibility to the effect of the feature estimated with β_h^1 .


This averaging, however, comes at a cost. Collapsing each country's trajectory into a single time series mean imposes an extreme smoothing on the degree of heterogeneity in the panel, which may at times be considerably larger than suggested by the country averages. Consequently, the gap between the impulse responses evaluated at the lower and upper ends of the distribution is highly conservative, and a statistically significant difference between the two states is therefore difficult to attain (as shown by the right hand side charts associated with specification 2 on Figures 8b, 9b and 10b). The two approaches are complementary: consistent findings across both specifications provide stronger evidence that the structural characteristic of interest influences monetary policy transmission in the euro area. This is indeed what we observe for the different economic features examined below.

4.2 Baseline results

4.2.1 The role of the debt characteristics

Fixed versus variable interest rates (Figure 8). We first analyze the role of a variable versus fixed interest rate regime. Although Italy, Portugal and Spain have undergone a regime shift over time from predominant variable rates to fixed rates, Germany, France and the Netherlands all have a low share of variable rates loans while Luxembourg, Finland, Greece, Lithuania and Latvia all have a high one. Specification 1 reveals that a high share of variable rate loans strengthens the transmission of monetary policy to output and inflation, confirming the importance of the exposure to interest rate changes for the transmission to aggregates. Interest rate changes pass through more rapidly and directly to borrowers' debt servicing costs, amplifying the effect of the shock. Indeed, a rise in the policy rate that translates immediately into higher debt servicing costs for variable-rate borrowers, reduces households' disposable income and weighs on firms' debt servicing capacity and investment resources, thereby amplifying the contractionary effect on aggregate demand. The subsequent disinflationary effect on prices then follows as the demand contraction propagates through the economy. Fixed interest rates insulate borrowers from the immediate repricing of their debt and delays the pass-through, as a shock absorber. Specification 2 corroborates these findings. The interaction coefficient between the shock and the share of variable rate loans is negative and significant for both GDP and HICP across most horizons, indicating that a higher level of the characteristic amplifies the contractionary effect of the shock. The

Figure 8: The role of fixed versus variable interest rate loans to firms and households




Notes: The high versus low classification is based on the pooled median ratio, equal to 85,0% (minimum 21,4%, maximum 99,8%). The shaded areas show the 1-std confidence interval.

Wald test					
H	p-value	H	p-value	H	p-value
H0	0.07	H7	0.01	H14	0.12
H1	0.20	H8	0.01	H15	0.09
H2	0.01	H9	0.01	H16	0.29
H3	0.05	H10	0.03	H17	0.51
H4	0.02	H11	0.01	H18	0.47
H5	0.04	H12	0.00		
H6	0.00	H13	0.01		

Wald test					
H	p-value	H	p-value	H	p-value
H0	0.43	H7	0.12	H14	0.17
H1	0.95	H8	0.08	H15	0.13
H2	0.56	H9	0.03	H16	0.42
H3	0.11	H10	0.04	H17	0.56
H4	0.13	H11	0.03	H18	0.26
H5	0.07	H12	0.03		
H6	0.14	H13	0.11		

(b) Specification 2



Notes: The left charts report the effect of a one-std monetary policy shock when the characteristic is at its sample mean; the middle charts report the marginal effect associated with a 1 std increase in the characteristic above its mean. The right charts display the total effect of the shock, evaluated at the 25th and 75th percentiles of the cross-country distribution of the characteristic, where each country's value is averaged over the sample period. This averaging is conservative, as it smooths out within-country variation over time and therefore narrows the gap between the two evaluation points.


total effect evaluated at z_{p25} and z_{p75} — corresponding to countries at the lower and upper ends of the distribution of country-average values of the characteristic — confirms that euro area countries with a higher share of variable rate loans experience a larger contraction in both GDP and HICP, though the gap between the two evaluation points is more conservative, as expected given the smoothing induced by the country-level averaging over the sample.

The maturity of loans (Figure 9). We now study the role of a high share of long-term loans proxied by the share of loans to firms and households with a maturity equal to five years or more. Germany, Finland, France and the Netherlands consistently display a high proportion of long-term loans while Ireland, Italy, Greece and Luxembourg systematically exhibit a low share. A higher share of long-term loans can affect monetary policy transmission in two opposing ways. On the one hand, it mutes the price of credit channel. The long-term share matters here through the pricing horizon: when loans are predominantly granted at long maturities, lenders price them off the corresponding long end of the yield curve rather than off the short-term policy rate. Because the long end responds only partially to a change in the policy rate a rise in the policy rate passes through only weakly to the rates charged on new long-term loans. The dominance of long maturities thus mutes the pass-through (even independently of whether the rate is fixed or variable). Thus, average rates on new loans tend to react less to a monetary policy shock in countries where long-term mortgages dominate. By shielding firms and households from a strong repricing of the cost of debt, a high share of long-term loans thus reduces the sensitivity of the non-financial private sector to policy rate changes, acting as a shock absorber. On the other hand, it can amplify the bank-capital channel through the maturity mismatch on banks' balance sheets: a bank funding long-maturity assets with short-maturity liabilities sees its funding costs

rise immediately when rates increase, while the yield on its existing assets turns over only slowly. The resulting compression of net interest margins erodes net worth and leads capital-constrained banks to cut lending. The two effects work in opposite directions, and our results suggest that the muting via the price of credit channel dominates, though with an asymmetry between HICP and GDP. For HICP, Specification 1 reveals a clear and persistent divergence between states: a high share is correlated with a more muted and delayed adjustment. For GDP, results point in the same direction but hardly reach statistical significance. Specification 2 corroborates these findings: the interaction coefficient is positive and significant for HICP across most horizons, confirming the attenuation effect, while it remains close to zero and insignificant for GDP. The total effect IRFs evaluated at z_{p25} and z_{p75} show a persistent gap for HICP, validating the role of the maturity structure for monetary policy transmission for the average level of heterogeneity in the euro area.

Figure 9: The role of the maturity of loans to firms and households

(a) Specification 1




Notes: The high versus low classification is based on the median share of long-term loans to firms and households, equal to 66,3% (minimum 26,7%, maximum 82,6%). The shaded areas show the 1-std confidence interval.

Wald test					
H	p-value	H	p-value	H	p-value
H0	0.92	H7	0.29	H14	0.64
H1	0.40	H8	0.29	H15	0.71
H2	0.42	H9	0.16	H16	0.91
H3	0.28	H10	0.40	H17	0.67
H4	0.21	H11	0.53	H18	0.98
H5	0.92	H12	0.29		
H6	0.28	H13	0.49		

Wald test					
H	p-value	H	p-value	H	p-value
H0	0.09	H7	0.01	H14	0.08
H1	0.02	H8	0.01	H15	0.07
H2	0.01	H9	0.00	H16	0.16
H3	0.00	H10	0.01	H17	0.12
H4	0.00	H11	0.03	H18	0.08
H5	0.02	H12	0.03		
H6	0.01	H13	0.08		

(b) Specification 2




Notes: The left charts report, the effect of a one-std monetary policy shock when the characteristics is at its sample mean; the middle charts report, the marginal effect associated with a 1 std increase in the characteristic above its mean. The right charts display the total effect of the shock, evaluated at the 25th and 75th percentiles of the cross-country distribution of the characteristic, where each country's value is averaged over the sample period. This averaging is conservative, as it smooths out within-country variation over time and therefore narrows the gap between the two evaluation points.

The level of households indebtedness (Figure 10). We now test the sensitivity of monetary policy transmission to the ratio of loans over gross disposable income, a proxy for the level of households indebtedness. Here, the sample is restricted to the founding members of the euro area except Luxembourg, as they are the only countries for which the relevant data series are available at a monthly frequency. As shown by the third chart, in Appendix E.1, there is a high euro area cross-country heterogeneity along the debt ratio of households and this heterogeneity is remarkably constant over time. Belgium, Ireland, the Netherlands consistently display a high indebtedness while Austria and Italy systematically exhibit a low one. We expect this dimension to matter for the revenue channel of monetary policy and to work as an amplifier of transmission to real activity. When the central bank raises policy rates, borrowers face higher debt servicing costs. In economies where households have a larger share of their income committed to debt service, this redistribution translates into a larger fall in net disposable income, a sharper contraction in consumption expenditures (Flodén et al. (2017)) and ultimately a more pronounced decline in activity. Since households tend to prioritize mortgage payments over other expenditures when facing financial distress, this reinforces the drop in consumption expenditures. The results of both specifications confirm the plausibility of such channel. With Specification 1, the difference between the high and low indebtedness states is immediate, large and statistically significant throughout the entire IRFs for GDP, while no significant difference emerges for HICP at medium and longer term horizons. Specification 2 results are in line, with a negative and significant interaction coefficient for GDP across all horizons and an interaction coefficient indistinguishable from zero for HICP at the medium-term horizons that matter. The total effect IRFs evaluated at z_{p25} and z_{p75} support the pattern. The asymmetry between activity and price responses is however pronounced for this

characteristic: household indebtedness amplifies the effect of a monetary policy tightening on activity, while leaving prices largely unaffected.

Figure 10: The role of the level of household indebtedness

(a) Specification 1




Notes: The high versus low classification is based on the median ratio which is equal to 92,9% (with a minimum of 30,3% and a maximum of 256,2%). The shaded areas show the 1-std confidence interval.

Wald test					
H	p-value	H	p-value	H	p-value
H0	0.45	H7	0.00	H14	0.00
H1	0.01	H8	0.00	H15	0.00
H2	0.02	H9	0.00	H16	0.00
H3	0.18	H10	0.00	H17	0.00
H4	0.00	H11	0.00	H18	0.00
H5	0.02	H12	0.01		
H6	0.00	H13	0.00		

Wald test					
H	p-value	H	p-value	H	p-value
H0	0.50	H7	0.64	H14	0.60
H1	0.79	H8	0.75	H15	0.55
H2	0.83	H9	0.98	H16	0.70
H3	0.54	H10	0.94	H17	0.69
H4	0.35	H11	0.77	H18	0.69
H5	0.55	H12	0.78		
H6	0.56	H13	0.69		

(b) Specification 2



Notes: The left charts report, the effect of a one-std monetary policy shock when the characteristics is at its sample mean; the middle charts report, the marginal effect associated with a 1 std increase in the characteristic above its mean. The right charts display the total effect of the shock, evaluated at the 25th and 75th percentiles of the cross-country distribution of the characteristic, where each country's value is averaged over the sample period. This averaging is conservative, as it smooths out within-country variation over time and therefore narrows the gap between the two evaluation points.


4.2.2 The role of the sectoral specialization

Finally, we test whether a larger ratio of manufacturing-to-services value added influences monetary policy transmission (**Figure 11**). Austria, Estonia, Finland, Germany, Ireland, Italy, Lithuania, Slovakia, Slovenia consistently display a higher ratio of manufacturing-to-services in the value-added over the whole sample while France, Cyprus, Greece, Luxembourg, Malta, Latvia, the Netherlands systematically exhibit a lower one. The results of both specifications are aligned. In the short-run (only), a tightening monetary policy shock transmits more forcefully to the real activity when the share of the manufacturing sector relative to the services sector is higher, confirming that interest rate hikes have a stronger contractionary impact on GDP in euro area economies such as Germany (Datsenko and Fleck (2024)). The underlying mechanism relates to the interest rate / cost-of-capital channel (Georgiadis (2015), Hauptmeier et al. (2025)). The manufacturing sector is generally more sensitive to changes in interest rates due to several structural factors. First, it is more capital-intensive, relying heavily on external financing to finance investment (e.g. machinery, equipment). By contrast, the services sector, more labor-intensive, typically requires fewer large capital expenditures.

Second, manufacturing firms tend to hold significant inventories and operate with higher levels of debt, further amplifying their exposure to interest rate fluctuations. Third, manufacturing firms produce

Figure 11: The role of the sectoral composition

(a) Specification 1




Notes: The high versus low classification is based on the pooled median ratio of manufacturing-to-services value added, equal to 21% (minimum 4%, maximum 72%). The shaded areas show the 1-std confidence interval.

Wald test					
H	p-value	H	p-value	H	p-value
H0	0.36	H7	0.17	H14	0.37
H1	0.65	H8	0.09	H15	0.26
H2	0.23	H9	0.17	H16	0.20
H3	0.53	H10	0.09	H17	0.78
H4	0.88	H11	0.24	H18	0.93
H5	0.74	H12	0.10		
H6	0.06	H13	0.09		

Wald test					
H	p-value	H	p-value	H	p-value
H0	0.46	H7	0.30	H14	0.73
H1	0.40	H8	0.28	H15	0.66
H2	0.18	H9	0.19	H16	0.89
H3	0.20	H10	0.23	H17	0.76
H4	0.13	H11	0.26	H18	0.97
H5	0.07	H12	0.47		
H6	0.16	H13	0.87		

(b) Specification 2



Notes: The left charts report, the effect of a one-std monetary policy shock when the characteristic is at its sample mean; the middle charts report, the marginal effect associated with a 1 std increase in the characteristic above its mean. The right charts display the total effect of the shock, evaluated at the 25th and 75th percentiles of the cross-country distribution of the characteristic, where each country's value is averaged over the sample period. This averaging is conservative, as it smooths out within-country variation over time and therefore narrows the gap between the two evaluation points.

more durable goods, including investment goods (Carlino and Defina (1998), Peersman and Smets (2002)). The demand for these goods is more sensitive to interest rates because they are often financed through borrowing and their consumption can be postponed, exhibiting high intertemporal elasticity of substitution. Therefore, economies with a larger manufacturing sector experience stronger and more immediate output effects from monetary policy shocks compared to services dominated economies.

The stronger effect of a contractionary monetary policy shock in economies with a larger manufacturing share suggests that interest rate sensitivity is the dominant transmission channel, rather than differences in price rigidity across sectors. Although the higher price stickiness in the service sector³³ could in principle generate stronger responses in services than in manufacturing, our results indicate that the capital-intensive nature of manufacturing activities dominates. We also observe that the significant difference between the two IRFs for activity does not translate into a significant difference for HICP.


To complement the previous analysis, we compare the euro area output responses of the manufacturing and services sectors to a tightening monetary policy shock using local projections. Figure 12 illustrates how sectoral characteristics contribute to differences in monetary policy transmission, on top of country-level heterogeneities. Because monthly industrial production data for the services sector are available only from January 2015 onward, we construct a longer series by interpolating the quarterly services GDP series³⁴, using the Chow and Lin (1971) method, with the help of the monthly euro area indicator

³³Services tend to be more labor-intensive, with wages set through contracts that adjust infrequently. Moreover, several service industries (health or education) are subject to regulatory constraints and, sometimes, administered prices that further slow price adjustments relative to manufacturing (Allayioti et al. (2024)).

³⁴The GDP of the service sector is the sum of wholesale and retail trade, transport, accommodation and food service activities, information and communication, financial and insurance activities, real estate activities, professional, scientific and technical activities, public administration, defense, human health, education, social work activities, art, entertainment and recreation GDP.

for total industrial production. The manufacturing sector exhibits a markedly stronger and more persistent decline following a monetary policy shock. Its response becomes significantly negative within a few months and reaches a trough of around -0.8% before gradually stabilizing, consistent with its higher capital intensity, stronger dependence on external financing and higher durability of goods produced. By contrast, the services sector reacts more modestly, with a smaller and smoother decline that remains close to -0.3% throughout the horizon. Overall, these results indicate substantial sectoral heterogeneity in monetary policy transmission, a dimension that may play an important role in shaping the transmission of monetary policy alongside cross-countries heterogeneities.

Figure 12: From country heterogeneity to sectoral heterogeneity



Notes: The figure reports the euro area industrial production responses of the manufacturing and services sectors to a one-standard-deviation monetary policy shock, estimated with local projections. The horizontal axis denotes the horizon in months. Shaded areas depict the one-standard-deviation confidence bands.

4.3 Evidence of offsetting effects across country characteristics

The three debt characteristics (maturity profile of loans, proportion of variable versus fixed rate, degree of households indebtedness) might not operate independently. Their respective influences on the transmission of monetary policy may interact and amplify or offset one another. To test this mechanism, we augment our baseline specification by simultaneously interacting the monetary policy shock with two structural characteristics, $\tilde{Z}_1$ and $\tilde{Z}_2$, as follows:

$$\begin{aligned}
y_{t+h} = & \alpha_h + \beta_h^1 \cdot MPS_t + \beta_h^2 \cdot \tilde{Z}_{1,t-1} + \beta_h^3 \cdot \tilde{Z}_{2,t-1} \\
& + \beta_h^4 \cdot MPS_t \times \tilde{Z}_{1,t-1} + \beta_h^5 \cdot MPS_t \times \tilde{Z}_{2,t-1} \\
& + \beta_h^6 \cdot MPS_t \times \tilde{Z}_{1,t-1} \times \tilde{Z}_{2,t-1} + \sum_{k=1}^p \phi_k^h \cdot X_{t-k} + v_{t+h}
\end{aligned} \tag{4.3.1}$$

where the triple interaction coefficient, β_6 , determines how the sensitivity of monetary policy transmission to $\tilde{Z}_{1,t-1}$ depends on the level of $\tilde{Z}_{2,t-1}$ (and vice versa), such that the marginal effect of $\tilde{Z}_{1,t-1}$ on the transmission of the monetary policy shock is:

$$\frac{\partial^2 y_{t+h}}{\partial MPS_t \partial \tilde{Z}_{1,t-1}} = \beta_h^4 + \beta_h^6 \cdot \tilde{Z}_{2,t-1} \quad (4.3.2)$$

and the total effect of the monetary policy shock, conditional on $\tilde{Z}_{1,t-1}$ and $\tilde{Z}_{2,t-1}$ is:

$$\frac{\partial y_{t+h}}{\partial MPS_t} = \beta_h^1 + (\beta_h^4 + \beta_h^6 \times \tilde{Z}_{2,t-1}) \times \tilde{Z}_{1,t-1} + \beta_h^5 \times \tilde{Z}_{2,t-1} \quad (4.3.3)$$


Interaction between debt maturity and interest rate regime (Figure 13). Taken separately, a high share of long-term loans attenuates monetary policy transmission, by attenuating the response of the cost of new credit (priced off the long end of the yield curve) while a high share of variable-rate loans amplifies it, by governing the effective rate paid on the stock of debt. In countries where both characteristics coexist, they may therefore exert offsetting effects on monetary policy transmission. The left panels show that a one-standard deviation contractionary monetary policy shock is associated with a statistically significant decline of approximately -0.5% in GDP and -0.3% in HICP when the share of long-term loans and the share of variable-rate loans are both at their euro area averages. The middle panels confirm the attenuation effect of long-term debt maturity, as a one-standard deviation increase in the share of long-term loans above its average reduces the contractionary effect of the monetary policy shock on both GDP and HICP. The right panels display the triple interaction coefficient β_h^6 , which captures the additional effect on transmission when both the share of long-term debt and the share of variable rates are simultaneously one standard deviation above their respective euro area averages. The coefficient is negative and significant across most horizons for HICP, and broadly negative though less precisely estimated for GDP. The negative signs indicate that the coexistence of a high share of variable-rate loans with a high share of long-term loans partially undoes the attenuation effect of long-term maturity: variable-rate exposure counteracts the insulation provided by longer debt maturities, as borrowers remain sensitive to interest rate repricing despite holding longer-term debt.

To assess whether the effects of these two characteristics on monetary policy transmission fully or only partially cancel out, we have conducted a Wald test of $H_0 : \beta_h^4 + \beta_h^5 + \beta_h^6 = 0$.³⁵ The test results are presented in Appendix E.3 using both clustered and DK standard errors. Overall, the two macroeconomic variables give a contrasting picture. For real GDP, the p-values are high across almost every horizon and under both standard error specifications, so that we fail to reject H_0 at any of the usual conventional levels, except at the 68% level for a handful of medium-term horizons

³⁵When both $\tilde{Z}_{1,t-1}$ and $\tilde{Z}_{2,t-1}$ are one standard deviation above their respective means ($\tilde{Z}_{1,t-1} = 1$ and $\tilde{Z}_{2,t-1} = 1$), the total effect of the monetary policy shock becomes $\beta_h^1 + \beta_h^4 + \beta_h^5 + \beta_h^6$. So the offsetting is total if $\beta_h^4 + \beta_h^5 + \beta_h^6 = 0$.

(p-value < 0.32). For HICP, the picture is more nuanced. At the shorter horizons we again fail to reject the hypothesis of complete offsetting at conventional levels, but the p-values decline markedly from horizon 6 onward, regardless of the standard errors used. The test results point to a rejection of H_0 at conventional levels, suggesting only a partial cancellation of the two effects in the long-run. These results underscore the importance of considering the joint structure of the non-financial private sector debt, rather than each dimension in isolation, when assessing the implications of cross-country heterogeneities for monetary policy transmission.

Figure 13: Interaction between debt maturity and interest rate regime.




Notes: The figure reports results from specification (4.3.1), in which the monetary policy shock is interacted simultaneously with the share of long-term loans ($\tilde{Z}_1$) and the share of variable rate loans ($\tilde{Z}_2$), both standardized and lagged by one period. The left panel shows β_1^h , the effect of a one-standard-deviation monetary policy shock when both characteristics are at their euro area averages. The middle panel shows β_4^h , the marginal attenuation from long-term maturity. The right panel shows the triple-interaction coefficient β_6^h , capturing the additional effect on transmission when both characteristics are simultaneously one standard deviation above their averages. The horizontal axis denotes the horizon in months. Shaded areas depict the 68% confidence bands.

Interaction between debt maturity and level of households indebtedness³⁶ (Figure 14). Taken in isolation, higher household indebtedness tends to amplify the effects of monetary policy, as more leveraged households are sensitive to changes in borrowing conditions. By contrast, a larger share of long-term debt attenuates transmission by slowing the pass-through of interest rate changes. When these two features coexist, their effects are therefore expected to partially offset each other. The left panels show that a one-standard deviation contractionary monetary policy shock is associated with a statistically significant decline of approximately -0.3% in GDP and -0.2% in HICP when the share of long-term loans and of household loans-to-income ratio are both at their euro area averages. The middle panels confirm the amplifying effect of a high level of household indebtedness, as a one-standard deviation increase in the household loans-to-income ratio above its average is negative and, consequently, reinforces the contractionary effect of the monetary policy shock on both GDP

³⁶The perimeter of the two characteristics is aligned in this specification. As the analysis focuses on household indebtedness, the share of variable-rate loans is defined over loans to households only, rather than over loans to both households and firms.

and HICP. The right panels display the triple interaction coefficient β_6^h , which captures the additional effect on transmission when both the share of long-term debt and the level of household indebtedness are simultaneously one standard deviation above their respective euro area averages. The estimated coefficient associated with this triple interaction term is both positive and statistically significant for both HICP and GDP, showing that an increase in the share of long-term loans mitigates the amplifying effect associated with an increase in the degree of households indebtedness. Although countries with highly indebted households remain more exposed to monetary tightening, longer debt maturities dampen this sensitivity by slowing down the pass-through of interest rate changes to the price of credit.

Figure 14: Interaction between debt maturity and the level of HH indebtedness




Notes: The figure reports results from specification (4.3.1), in which the monetary policy shock is interacted simultaneously with the household loans-to-income ratio ($\tilde{Z}_1$) and the share of long-term loans ($\tilde{Z}_2$), both standardized and lagged by one period. The left panel shows β_1^h , the effect of a one-standard-deviation monetary policy shock when both characteristics are at their euro area averages. The middle panel shows β_4^h , the marginal amplification from a high household loans-to-income ratio. The right panel shows the triple-interaction coefficient β_6^h , capturing the additional effect on transmission when both characteristics are simultaneously one standard deviation above their averages. The horizontal axis denotes the horizon in months. Shaded areas depict the 68% confidence bands.

For real GDP, the test results shown in Appendix E.3 are robustly consistent with full cancellation of the two effects: we fail to reject H_0 at conventional levels across almost every horizon and under both specifications of the standard errors. For HICP, the conclusions are more nuanced since they depend on the standard errors used. Under clustered standard errors, we reject H_0 strongly across the short and medium horizons — and even at the longer horizons once a 68% confidence level is considered — pointing to only a partial cancellation. Under DK standard errors, the p-values are higher, yet they remain below the 0.32 threshold for most of the short and medium horizons, so that we cannot rule out a rejection of the complete offsetting hypothesis, at least for short and medium term horizons.

Interaction between interest rate regime and level of households indebtedness³⁷ (Figure 15). Finally, taken separately, a high share of fixed-rate loans dampens monetary policy transmission, while a high level of household indebtedness amplifies it. The left panels show that a one-standard deviation contractionary monetary policy shock is associated with a statistically significant decline of approximately -0.4% in GDP and -0.2% in HICP when the share of fixed rate loans and the share of household loans-to-income ratio are both at their euro area averages. The middle panels confirm that an increase in the level of households debt reinforces the contractionary effect of the monetary policy shock, as a one-standard deviation increase in the household loans-to-income ratio above its average has a negative effect on GDP and on HICP (after horizon 8). The right panels display the triple interaction coefficient β_6^h , which captures the additional effect (beyond the effects of each characteristic taken individually) on transmission when both the share of fixed rate loans and the level of household indebtedness are simultaneously one standard deviation above their respective euro area averages. The results presented

Figure 15: Interaction between interest rate regime and the level of HH indebtedness



Notes: The figure reports results from specification (4.3.1), in which the monetary policy shock is interacted simultaneously with the household loans-to-income ratio (Z_1) and the share of fixed rate loans (Z_2), both standardized and lagged by one period. The left panels show β_1^h , the effect of a one-standard-deviation monetary policy shock when both characteristics are at their euro area averages. The middle panel shows β_4^h , the marginal amplification from a higher household loans-to-income ratio. The right panel shows the triple-interaction coefficient β_6^h , capturing the additional effect on transmission when both characteristics are simultaneously one standard deviation above their averages. The horizontal axis denotes the horizon in months. Vertical bars depict the 68% confidence intervals.

in the right panels are consistent with the offsetting mechanism expected to arise from the interaction between these two characteristics. Indeed, the estimates of β_6^h are positive and statistically significant across most horizons for HICP, and to a lesser extent, for GDP, indicating that a higher share of fixed-rate loans attenuates the amplification effect associated with elevated household indebtedness. In other words, this analysis confirms that when a larger fraction of debt is insulated from contemporaneous

³⁷The perimeters of the two characteristics are aligned in this specification. As the analysis focuses on household indebtedness, the share of variable-rate loans is defined over loans to households only, rather than over loans to both households and firms.

interest rate fluctuations (as the protection afforded by fixed-rate contracts limits the adjustment in debt servicing costs), the effect of a higher degree of households indebtedness on monetary policy transmission is reduced.

For real GDP, the test results in Appendix E.3 are again robustly consistent with full cancellation: we fail to reject H_0 at conventional levels across almost every horizon and under both specifications of standard errors. For HICP, on the contrary, the null hypothesis of full offsetting is rejected at most horizons under both specifications: most p-values even fall below 0.05 with clustered standard errors, while under DK they are higher but almost always below 0.32 and close to 0.10 for most horizons.

Thus, for the three key sources of cross-country heterogeneity in the euro area examined pairwise in this subsection, their combined effect on monetary policy transmission is smaller in absolute terms than the sum of their individual effects. The illustration of these offsetting mechanisms raises a broader question: does the partial cancellation between amplifying and attenuating characteristics reflect a fortunate coincidence, or does it point to a deeper endogenous regularity in how national economic structures are formed? The offsetting need not be coincidental, nor confined to debt characteristics. Several of the characteristics that push transmission in opposite directions are plausibly related, either responding to common underlying forces or being chosen, at least in part, to compensate for one another. A high share of variable-rate loans, which amplifies transmission, may for instance coexist with longer maturities or lower household leverage precisely because borrowers, lenders and regulators adjust debt contracts in response to the interest-rate risk that variable-rate exposure creates. The same logic could extend to other dimensions of heterogeneity: an economy with a larger manufacturing sector and therefore a stronger exposure to the cost-of-capital channel might also be more likely to feature other characteristics that curb its exposure to interest rate changes, precisely because firms and banks adapt financing structures to the interest-rate risk that a capital-intensive financed sector entails. If so, structural characteristics are not drawn independently across countries, and the low heterogeneity of MPT at the aggregate level would reflect an endogenous, self-stabilizing configuration of national characteristics, rather than the fortuitous cancellation of unrelated features.³⁸

Overall, this analysis underscores the importance of integrating both micro- and macro-level approaches when studying monetary policy transmission heterogeneity. Reconciling these approaches provides a more comprehensive understanding of how heterogeneities shape the effectiveness and distribution of monetary policy outcomes across euro area countries and helps understand how key sources of heterogeneity for monetary policy transmission can coexist with a contained heterogeneity of countries' responses.

³⁸One way to investigate this conjecture would be to test whether amplifying and attenuating characteristics systematically co-occur across countries, an exercise that we keep for future research.

5 Assessing the evolution of cross-countries heterogeneity over time: a FAVAR approach

In this section, we now assess the evolution of cross-country heterogeneity in monetary policy transmission between January 2000 and July 2025. Over this period, euro area countries have experienced important shocks and structural shifts. The macroeconomic environment evolved from the Great Recession to the post-pandemic inflation surge, further intensified by the energy price shocks that followed the outbreak of the war in Ukraine. The financial landscape was transformed by the Global Financial Crisis and the euro area sovereign debt crisis. In these contexts, the monetary policy framework itself evolved from conventional interest rate instruments to an extensive toolkit of unconventional measures. These shocks and transformations may have modified the magnitude and the channels of monetary policy transmission for both the euro area and its individual member states. Moreover, the 2022-2025 monetary policy cycle offers a unique opportunity to evaluate whether the degree of heterogeneity in monetary policy transmission, and with it the degree of fragmentation in the monetary union, has changed over time.

The standard local projections used until now are well known to be unstable across sample length and composition. Against this background, the FAVAR model delivers estimates that are robust to sample length and composition, allowing us to conduct a rolling-window analysis. In addition, the FAVAR identifies monetary policy shocks within a rich information set that approximates the wide range of variables the central bank can monitor. This high data dimensionality reduces the risk of omitted variable bias and makes the model particularly suitable for estimation over short subsamples, as the breadth of the cross-section offsets the shorter time dimension inherent to a rolling-window analysis. Together, these two features are essential for analyzing the evolution of monetary policy transmission in a rolling-window framework.

5.1 The FAVAR model

5.1.1 Data

Our monthly FAVAR is estimated on a cross-country dataset covering the EA as well as all its founding members except Ireland³⁹: France, Germany, Italy, Spain, Portugal, Austria, the Netherlands, Finland, Belgium and Luxembourg. The construction of this panel dataset, bringing together a large set of macroeconomic, monetary, financial and banking indicators is a key contribution of this paper. For each country and for the euro area, we include eleven variables: real GDP, retail trade turnover, unemployment rate, HICP all items, HICP food, HICP energy, producer price index, stock market

³⁹Ireland is excluded because of the atypical behavior of its GDP that may distort the identification of the factors.

indices, the 10-year sovereign spread,⁴⁰ credit to households, as a proxy for credit to the private non-financial sector, and credit to non-monetary financial institutions, as a proxy for credit to the private financial sector.⁴¹ In addition, we also include, as global variables, the nominal effective exchange rate and oil prices.⁴² Finally, a short-term interest rate, used as a proxy for the policy rate, is included as the only directly observed variable in our VAR. The policy rate is measured using a 3-month Euribor *shadow rate* based on Golinski and Spencer (2024)).⁴³ The monthly frequency of our FAVAR model limits the inclusion of additional macroeconomic variables, such as labor market, productivity and investment indicators, that are commonly used in quarterly FAVAR estimations prevalent in the literature. In contrast, our focus is deliberately on banking and financial indicators, which are available and relevant at lower frequencies, but appear less frequently in the seminal contributions of the FAVAR literature. The monthly real GDP series is obtained by interpolating quarterly GDP with its corresponding industrial production index via the Chow and Lin (1971) method, as for the local projections presented previously.⁴⁴ To address stationarity issues, we use the year-on-year growth rate of all time series, except for the interest rate (in level) and for the unemployment rate (12-month first difference). Detailed variable definitions, data sources and transformations are provided in Appendix F. Thus, our monthly FAVAR is estimated using a balanced panel of 121 time series spanning from January 2000 to July 2025, excluding the COVID-19 period from March to December 2020, in line with the previous sections. Excluding this period is necessary in a factor-model setting, as factor identification is sensitive to structural breaks in the data.⁴⁵

Two elements are noteworthy. First, following Boivin et al. (2009), we include both aggregate euro area variables and country-specific series in the factor estimation stage of the FAVAR. While this choice may raise concerns about redundancy and potential bias in factor extraction, these concerns are alleviated by the robustness check that we have performed: re-estimating the model without EA variables yields virtually identical results, indicating that their inclusion does not affect the estimated factors or the identification of the implied common component. In this sense, EA variables do not contribute materially

⁴⁰The sovereign spread of country i is defined as its 10-year government bond yield minus the 10-year government bond yield of Germany, which serves as the benchmark. Consequently, for this variable, euro area and Germany are not included.

⁴¹We additionally considered a specification incorporating the growth rate of credit to NFCs alongside credit to households. Our results are robust to this extension, in both the full-sample and rolling-window exercises, since credit conditions are already well captured by the growth rates of mortgages and of credit to non-MFI sectors. We nonetheless retained the more parsimonious specification: adding the credit to NFCs series brings the number of credit-related series to 33 out of 121, introducing largely redundant information into the panel and generates occasional convergence difficulties in the Gibbs sampler over the short windows used in the rolling-window exercise. Thus, excluding these series yielded a better behaved specification without materially affecting the results.

⁴²Our baseline results remain robust to the inclusion of additional global controls, such as the corporate bond spread, monetary aggregates and US-related variables.

⁴³The shadow rate is derived from Banque de France internal calculations. It applies the methodology of Golinski and Spencer (2024), developed for the US, to the euro area.

⁴⁴Our tests showed that our results are not sensitive to the use of the interpolated monthly GDP series rather than using the monthly industrial production series.

⁴⁵An extension of this paper would consist in implementing the approach of Lenza and Primiceri (2022), in order to explicitly account for the heightened volatility observed during the COVID-19 period without excluding these data from the sample. However, since our focus is on structural analysis and not on forecasting, we chose, at this stage, to exclude this period.

to factor identification. Their inclusion is nevertheless useful from an interpretative standpoint. The FAVAR framework is designed to recover both country-specific and aggregate responses to a common monetary policy shock, making it particularly well suited to analyze heterogeneity in transmission across member states. Indeed, including euro area variables allows us to directly obtain euro area-wide estimates within the same framework and, consequently, to benchmark national dynamics against aggregate ones. Moreover, the presence of euro area variables provides a useful internal consistency check. In a well-specified model, common factors should load more strongly on euro area aggregate variables, while idiosyncratic components should play, by construction, a smaller role. This pattern is confirmed in our empirical results, further supporting the validity of the specification. Overall, retaining EA variables enhances interpretability at no cost in terms of factor estimation.

Second, in addition to household credit, we have chosen to include the credit granted by MFI (mostly banks) to non-monetary financial institutions. These institutions (investment funds, insurance corporations, pension funds and the other financial intermediaries) all operate between banks and ultimate borrowers or asset markets and constitute the core of the non-bank financial intermediation (NBFIs). Their inclusion is motivated by several considerations. First, they account for a quantitatively important share of total credit granted by banks, often exceeding that of households and non-financial corporations, with substantial cross-country heterogeneity. Second, while credit to NBFIs comoves with traditional bank lending over the cycle and therefore contributes, like any other credit proxies, to capturing general credit conditions, it also exhibits distinct dynamics in terms of timing, volatility and speed of adjustment that are informative for identifying monetary policy shocks via their associated transmission channels. Indeed, a monetary policy tightening that deteriorates the funding conditions of NBFIs, regardless of whether they rely on bank-based or market-based financing, compels them to deleverage under binding constraints. This balance sheet adjustment leads to asset sales and exerts downward pressure on asset prices. The decline in remaining asset valuations further reduces collateral values, reinforcing deleveraging pressures through a collateral spiral. Empirically, we find that, everything else equal, including this variable helps resolve the commonly observed *price puzzle* of asset prices response to a monetary policy shock, as it acts as a natural bridge between the credit and asset price blocks of the model.

5.1.2 Methodology

Our standard FAVAR is composed of a transition equation (VAR model) for the latent factors and a measurement equation linking these factors to the observed variables.

Transition equation:

$$\begin{pmatrix} F_t \\ R_t \end{pmatrix} = B(L) \begin{pmatrix} F_{t-1} \\ R_{t-1} \end{pmatrix} + v_t \quad (5.1.1)$$

This transition equation describes a VAR(p), with $p = 1$ ⁴⁶ and where F_t is a (Kx1) vector of unobserved factors and R_t is a shadow rate used as a proxy for the policy rate. R_t is assumed to be the sole factor observable to both the econometrician and the central bank. This choice is standard and corresponds to the "preferred specification" (Bernanke et al. (2005)) in most FAVAR models used for monetary policy analysis. The disturbance term v_t is assumed to follow an i.i.d. process with mean zero and diagonal covariance matrix Σ .

Measurement equation:

$$X_t = \Lambda^F F_t + \Lambda^Y R_t + e_t = \Lambda C_t + e_t \quad (5.1.2)$$

$$\begin{pmatrix} X_{1,t} \\ \vdots \\ \vdots \\ X_{N,t} \\ R_t \end{pmatrix} = \begin{pmatrix} \Lambda^{11} & \Lambda^{12} & \dots & \Lambda^{1K} & \Psi^{11} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \Lambda^{N1} & \Lambda^{N2} & \dots & \Lambda^{NK} & \Psi^{N1} \\ 0 & 0 & \dots & 0 & 1 \end{pmatrix} \begin{pmatrix} F_t^1 \\ F_t^2 \\ \vdots \\ F_t^K \\ R_t \end{pmatrix} + \begin{pmatrix} e_{1,t} \\ \vdots \\ \vdots \\ e_{N,t} \\ 0 \end{pmatrix}$$


This measurement equation decomposes the fluctuations in each series into a series-specific component e_t and a common euro area component, C_t consisting of K factors and the policy rate R_t . The factors are unobserved and assumed to capture most of the information contained in the dataset X_t . Thus, Λ is a (NxK) matrix of factor loadings that maps the factors used in the VAR to the original series of interest. The loading matrix Λ is defined such that the financial variables (stock prices, government bond spreads and exchange rate) are allowed to have a contemporaneous relationship with the nominal interest rate, i.e. non-zero elements of Ψ . These variables are categorized as *fast moving* as opposed to the other variables classified as *slow moving*.

Factors selection. Several approaches have been proposed in the literature to determine the appropriate number of factors to include in the VAR. One commonly used method relies on the three information criteria proposed by Bai and Ng (2002). In our case, these criteria converge slowly and suggest retaining a too large number of factors, with up to twelve factors, which would be impractical in a VAR setting. We therefore rely on the cumulative variance share criterion. The inspection of the scree plot, in Figure 16, which reports the share of variance explained by each successive factor (individually and cumulatively), suggests retaining five factors. Indeed, the scree plot exhibits a kink after the fifth factor

⁴⁶We chose $p=1$, as in Boivin et al. (2009). Each additional lag requires estimating a further set of VAR coefficients, which highly increased the computational cost in the rolling-window framework, where the model is re-estimated over a large number of windows.

and the first five ones explain a large and conventionally acceptable share of the total variance of the data (72%). This specification also yields impulse responses to the monetary policy shock that are fully consistent with standard priors and is in line with standard practice with factor models estimated for the euro area. Finally, beyond five, the marginal contribution of each additional factor becomes negligible, each adding little additional explanatory power (less than 5% of the total variance each). Appendix G.1 presents the top ten series that load the most heavily on each of the five factors. We study the sensitivity of our IRFs to the number of factors in Appendix G.2. Finally, given that our panel dataset X_t includes variables related to activity, prices, monetary, banking and financial conditions, the factors are intended to capture broad dimensions such as overall price levels, economic activity and monetary, financial and banking conditions. However, no restrictions are imposed on the loadings matrix to force an economic interpretation of the factors. Consequently, no factor can be uniquely associated with a single dimension, and, by construction, the factors do not have inherent economic significance and our results are not sensitive to alternative orderings of the factors in the VAR.

Figure 16: Factors selection via scree plot inspection



Notes: The scree plot reports the share of total variance of the data explained by each of the first 40 factors, ordered from largest to smallest. The share attributed to factor k is its eigenvalue divided by the sum of all eigenvalues (blue, individual); the cumulative curve (orange) reports the running total. The elbow of the individual curve, beyond which each additional factor adds little explanatory power, supports retaining five factors.

Monetary policy shock identification (Baumeister et al. (2013)). As implied by the transition equation, the policy rate is ordered last and the monetary policy shock is identified using a recursive ordering. With this Cholesky identification scheme, the monetary policy shock is the only disturbance that does not contemporaneously affect the latent factors within the month. The measurement equation (5.1.2) accounts for the contemporaneous correlations between the fast-moving variables and the policy rate via the factor loadings matrix which allows the fast moving variables, used to identify the factors, to adjust contemporaneously to the interest rate shock. The $\Lambda^Y R_t$ term allows the measurement

equation to isolate from the factor space the part of the factors that is explained by the interest rate today, and, consequently, to keep in the factor space only the component of the factor that is not affected contemporaneously by the movements in the monetary policy rate. This approach mitigates the disadvantages of a purely recursive identification scheme and legitimates a recursive ordering with the interest rate ordered last, after all the factors. Thus, this identification allows us to estimate impulse responses of both F_t and R_t to a one-standard deviation monetary policy shock and to finally recover IRFs for all variables in X_t .

5.1.3 Estimation

As in [Baumeister et al. \(2013\)](#) and [Mumtaz \(2010\)](#), the model is estimated, both on full sample and on successive windows, via the single-step Bayesian likelihood approach with Gibbs sampling algorithm to approximate the joint posterior distribution. This enables the simultaneous estimation of the factors and of the VAR dynamics.⁴⁷ First, an initial set of latent factors is extracted from the standardized panel data using principal component analysis to provide starting values for the sampler. Second, conditional on the current draw of the factors and the policy rate, the parameters of the observation equation are sampled under conjugate prior assumptions: factor loadings are assigned normal priors, while the idiosyncratic variances are assumed to follow inverse-gamma priors. Third, conditional on the factors and policy rate, the parameters (coefficients and the covariance matrix of the innovations) of the transition equation are estimated under a conjugate normal-inverse-Wishart prior. The coefficients of the VAR are assigned a normal Minnesota-type prior, while the covariance matrix of the innovations is assigned an inverse-Wishart prior. Fourth, the latent factors are updated by jointly drawing their entire time path using the [Carter and Kohn \(1994\)](#) algorithm. This algorithm consists of a forward path using the Kalman filter, which recursively computes the conditional distribution of the factors given information available up to each point in time (previous periods), followed by a backward simulation smoother that draws the full factor trajectory conditional on the complete sample (next periods). Together, these steps yield a full factor path consistent with all observed data. The resulting factors are then used in the next iteration of the Gibbs sampler, ensuring convergence to the posterior distribution of the model parameters. For further details, especially regarding the priors definition, we refer the reader to the technical appendix of [Baumeister et al. \(2013\)](#); the difference here being that we consider a FAVAR with constant parameters rather than a time-varying FAVAR. This procedure yields 1,000 effectively stored posterior draws.⁴⁸ Point estimates are then constructed as the median draw at each horizon, while credible intervals are obtained from the distribution of the retained draws

⁴⁷While more computationally demanding than the two-step principal component method ([Bernanke et al. \(2005\)](#)) or its variant ([Boivin et al. \(2009\)](#)), this approach offers the flexibility to accommodate a time-varying framework, a feature that could prove useful for future extensions with a time-varying FAVAR model.

⁴⁸Within this Bayesian framework, we chose 12,000 Gibbs sampling iterations, of which the first 2,000 are discarded as burn-in. To mitigate the serial correlation across posterior draws, we apply a thinning factor of 10, retaining one draw every ten iterations.


(16th–84th percentiles for the 68% interval and 5th–95th percentiles for the 90% one). This Bayesian approach allows us to compute the IRFs and all heterogeneity metrics with their own credible intervals.

5.2 Full-sample estimation: alternative measure of cross-country heterogeneity


In this section, we analyze the effects of a one-standard deviation monetary policy shock on all the country-specific variables included in our estimation. We do so, on Figure 17, by comparing, for all variables, the IRFs across countries. To provide an order of magnitude of the shock size, a one-standard deviation monetary policy shock corresponds to an increase, on impact, of 12 basis points in the 3-month Euribor-based shadow rate that we use as a proxy for the short-term policy rate.

Appendix H presents, for each variable and each country, the IRF with its 68% and 90% credible intervals. Overall, the results indicate a high degree of statistical significance, as almost all impulse responses, out of the 122 variables in total, are significant at the 90% confidence level, with only isolated exceptions. First, all variables exhibit the expected sign for nearly all countries. Second, a visual inspection of the IRFs across countries points to a contained degree of cross-country heterogeneity in the transmission of monetary policy for most variables. This visual impression is corroborated by the computation, in Table 2, of the coefficients of variation (CV), following the same approach previously applied to the local projection estimates.⁴⁹

Figure 17: FAVAR IRFs to a 1-standard deviation monetary policy shock.



⁴⁹ Here again, using the CV rather than the standard deviation facilitates meaningful comparisons of the degree of heterogeneity across variables, as it scales dispersion relative to the magnitude of the cross-countries average response.



Notes: The figure reports the country-specific impulse responses to a one-standard-deviation monetary policy shock, recovered from the FAVAR model. The shock corresponds to an increase of about 12 basis points, on impact, in the 3-month Euribor-based shadow rate used as a proxy for the policy rate, and is identified recursively with the shadow rate ordered last. The model is estimated at monthly frequency, on a balanced panel of 121 series covering the euro area and ten founding members over January 2000–July 2025, excluding the COVID-19 period. Responses are expressed in percentage points. For each country and variable, the solid lines report the median response across draws.

Our FAVAR results deliver a picture of cross-country heterogeneity strikingly close to the one obtained with local projections, despite the two frameworks resting on different assumptions: dispersion is lowest for real GDP and prices (except food prices), and more pronounced for credit volumes and the unemployment rate. That two different methodologies converge on similar conclusions reinforces our reading that the low degree of heterogeneity in the transmission to macroeconomic aggregates is a robust feature of the data rather than an artifact of a particular empirical strategy.

Stock price responses to a monetary policy shock display the smallest degree of dispersion across all variables, in relation to the high integration of euro area equity markets. Key macroeconomic aggregates - real GDP as well as price variables, particularly PPI, HICP energy and HICP all items - stand out for their particularly low dispersion in their responses to a monetary policy shock, confirming the previous results obtained with local projections. The greater dispersion in HICP food responses to a monetary policy shock likely reflects structural heterogeneity across euro area countries. When the central bank tightens policy, aggregate demand and wage pressures decline, but the extent to which these effects translate into lower consumer food prices depends on national characteristics. In particular, different degrees of competition in the food retail sector might lead to various cost absorption and pass-through (Ciapanna and Rondinelli (2014)). Unlike food prices, HICP energy responses exhibit a high degree of homogeneity across euro area countries, owing primarily to the fact that euro area monetary policy

transmits to energy prices through two main channels that are susceptible to operate: the demand channel and the exchange rate channel, which affects the domestic cost of imported energy (particularly oil and gas). Since these two channels can affect all energy types — and therefore all countries in a relatively uniform way, cross-country heterogeneity in energy price responses remains limited, despite considerable variation in energy mixes across member states.

Table 2: Coefficient of variation across variables

Variables	Median	68% CI	90% CI
Real GDP	0.399	[0.360, 0.450]	[0.334, 0.483]
Unemployment rate	0.748	[0.686, 0.816]	[0.648, 0.858]
Retail trade turnover	1.800	[1.236, 2.994]	[1.03, 5.862]
HICP all items	0.476	[0.357, 0.736]	[0.309, 1.444]
HICP energy	0.336	[0.306, 0.375]	[0.290, 0.414]
HICP food	2.243	[1.651, 4.254]	[1.345, 13.027]
PPI	0.282	[0.265, 0.312]	[0.257, 0.346]
Spread	1.051	[0.934, 1.246]	[0.863, 1.564]
Stock prices	0.194	[0.148, 0.411]	[0.132, 0.763]
Households credit	1.221	[0.982, 1.525]	[0.864, 1.828]
non-MFI credit	1.034	[0.981, 1.267]	[0.776, 1.552]

Notes: The table reports the coefficient of variation (CV) of the FAVAR impulse responses across the founding EA countries (except Ireland), by variable. This metric is computed at each draw of the Gibbs sampling; the table reports, over the 1,000 retained posterior draws, the median and the 16th–84th (68%) and 5th–95th (90%) percentiles. The CV is averaged across horizons before being summarized. Lower values, in bold, indicate the variables with the most homogeneous responses.

Moreover, the substantial cross-country heterogeneities in retail trade turnover responses to a monetary policy shock mirror the dispersion in household consumption responses. Several factors contribute to this higher heterogeneity of our monthly consumption proxy. As previously shown, euro area countries exhibit substantial differences in household debt levels and in the structure of their mortgage markets. In economies where variable rate mortgages predominate, an increase in interest rates immediately raises monthly mortgage repayments for a large share of households, directly reducing disposable income and, consequently, consumption expenditures. By contrast, in countries where fixed rate mortgages are more prevalent, the same rate hike has minimal short-term impact on existing borrowers' cash flows, leaving retail spending largely unaffected in the immediate term. In addition, the prevalence of consumer credit differs significantly across countries reinforcing the marked cross-country differences in households' exposure to interest rate fluctuations. Finally, monetary tightening also reduces asset values, including both financial assets and housing. The resulting wealth effects on consumption depend critically on the distribution and ownership of these assets among households.

Then, consistent also with previous results, the year-on-year growth rate of household credit displays a pronounced cross-country dispersion in its response to a monetary policy shock. A similar degree of heterogeneity characterizes the growth rate of credit to NBFIs granted by MFIs. Two forces are likely to shape the credit response to a monetary policy shock: the interest-rate sensitivity of the entities being

funded, on the demand side, and the balance-sheet health of the lending banks, on the supply side. Both may operate differently across countries. On the demand side, the NBFIs sector is heterogeneous across the euro area (investment funds dominate in Luxembourg, France and Germany, pension funds in the Netherlands, insurance corporations in France, Germany and Italy), and these segments are likely to react differently to monetary policy (Anyfantaki et al. (2026)). On the supply side, the balance-sheet capacity and the institutional and regulatory constraints of domestic banks also vary widely. The same shock therefore reaches credit to national NBFIs sectors with different intensities contributing to the dispersion observed.

The degree of heterogeneity in the unemployment rate responses, while higher than that observed for other macroeconomic aggregates, remains moderate and fully consistent with the level of heterogeneity estimated based on local projections. Finally, although sizable and higher than for most other variables, the degree of heterogeneity in spread responses can be considered relatively contained in view of the sovereign debt crisis that falls within our full-sample period.⁵⁰ Several factors may account for it, including differences in countries' fiscal positions when the shock hits, the maturity structure of outstanding sovereign debt, the composition of the investor base (domestic versus foreign holders), the size and liquidity of national bond markets and the financial health of banking sectors holding large quantities of domestic sovereign debt (*doom-loop* risks). But, overall, this result suggests that financial markets still perceive monetary tightening as a common macroeconomic shock rather than as a trigger for country-specific solvency concerns, pointing to little sign of financial fragmentation across euro area countries over the period.

Thus, for several variables, most notably real GDP and prices, we document lower cross-country heterogeneity than commonly reported in earlier studies. A candidate explanation we explore is the sample period: our estimation runs through 2025 and thus incorporates the 2022–2024 tightening cycle, whereas much existing evidence rests on samples ending at least a decade earlier. Since our FAVAR approach is close to the dynamic factor model of Corsetti et al. (2022), we isolate the sample effect by re-estimating it on their sample ending in 2016, holding everything else unchanged (same variables, iterations, draw selection and identification). On this otherwise-identical sample, heterogeneity is clearly higher for most variables, while the ranking across variables is broadly unchanged, pointing to the longer sample as at least one important source of our lower measured heterogeneity. Appendix I displays a comparative table of the coefficient of variation over 2000–2025 and 2000–2016, confirming that heterogeneity is systematically higher on the shorter pre-2016 sample. Two complementary factors may influence this pattern. First, the extended sample includes the effective lower bound and unconventional monetary policy period. Second, this period coincides with deeper banking integration:

⁵⁰This higher but relatively contained level of heterogeneity observed for the response of spreads should nonetheless be interpreted in light of the exclusion of Ireland and Greece from the sample.

the various reforms related to the Banking Union may have contributed to some convergence of MPT. Overall it is reassuring that our approach is able to replicate the findings of [Corsetti et al. \(2022\)](#) when using the same sample as they do.

5.3 Rolling windows

5.3.1 Evolution of heterogeneity over time

To assess the time variation in cross-country heterogeneity of monetary policy transmission, we implement a rolling window estimation designed to track the evolution of the dispersion in country-specific impulse responses to a monetary policy shock over time. To do so, our preferred metric is the cross-country standard deviation of IRFs, averaged over horizons. Specifically, we estimate the model over rolling windows of 12 years. This choice reflects a trade-off between, on one hand, having a sufficiently large number of windows—including an initial window with the minimum number of months to capture the smallest baseline set of events (favoring shorter windows)—and, on the other hand, ensuring reliable identification of monetary policy shocks within each subsample (favoring longer windows).⁵¹ For robustness checks, Appendix [J.3](#) reports results based on shorter 11-year windows, which yield highly comparable patterns for the evolution of heterogeneity. Thus, the first window spans the period from January 2000 to December 2013, which, given that the variables are expressed in year-on-year growth rates, effectively corresponds to the period 2001-2013. The window is then shifted forward by one month at a time, yielding a total of 140 overlapping windows, with the final window covering the period from August 2012 to July 2025.

For each window,⁵² we recover IRFs to a monetary policy shock for all countries and we measure cross-country heterogeneity as the standard deviation of the responses across countries. Because this dispersion metric is calculated at each horizon, we summarize it by averaging the standard deviations over horizons 6 to 24, capturing the short- to medium-term effects of monetary policy. Our results are also fully robust to averaging over the full 1–40 month horizon.


Importantly, this heterogeneity metric, like all others derived from the FAVAR framework, is computed within a Bayesian setting and therefore re-estimated at each iteration of the Gibbs sampler. This procedure allows us to construct uncertainty bands for our metric. Specifically, the 1,000 retained posterior draws provide a full posterior distribution for our metrics. We report the median as our point estimate, while the 68% and 90% credible intervals are obtained from the corresponding posterior percentiles. This procedure yields, for each variable, a time series of dispersion that captures the

⁵¹The 12-year window represents the shortest length that preserves stable estimation and credible identification of the monetary policy shock across the overwhelming majority of subsamples.

⁵²The rolling-window estimation requires 12,000 Gibbs sampling iterations for each of the 140 windows, which generates a substantial computational burden. To render the exercise feasible, we parallelized the code.

evolution of heterogeneity in monetary policy transmission, as illustrated in Figure 18. Given the window length, attributing individual peaks to specific shocks has to be done cautiously. However, it allows us to assess the plausibility of two non-mutually exclusive interpretations, each carrying distinct policy implications.

Figure 18: Evolution of cross-country heterogeneity in monetary policy transmission



Notes: Each panel plots, for a given variable, the evolution over time of cross-country heterogeneity in the transmission of a monetary policy shock, estimated from a FAVAR over rolling windows. For each window, heterogeneity is measured as the cross-country standard deviation of the impulse responses, averaged across horizons. The two lines differ only in the horizon range used for this averaging: H1–H40 (blue) and H6–H24 (grey). The horizontal axis reports the start and end dates of each rolling window; shaded bands depict the 68% and 90% credible intervals. Vertical dashed lines mark key monetary policy events.

The first is a *structurally driven heterogeneity hypothesis*. Under this view, structural reforms or major macroeconomic and financial disruptions, by durably altering the underlying structural features of the economies, would persistently shift the level of heterogeneity and generate a ratcheting-up (or down) pattern of dispersion over time, rather than transitory fluctuations around a stable level. A rise in such structurally driven heterogeneity would raise concerns for the conduct of a single monetary policy.

The second interpretation is a *cyclically driven MPT heterogeneity hypothesis*. Under this hypothesis, heterogeneity in monetary policy transmission would spike during some precise events before falling back to a low baseline level. This second form of heterogeneity is less worrisome for the central bank. It can even be searched for if during stress episodes, asymmetric MPT helps mitigate the uneven effects of certain shocks on countries. For instance, unconventional measures such as the launch of the Asset Purchase Programme (APP), aiming at restoring a smooth monetary policy transmission across all countries, or episodes of rapid and strong tightening might be associated with temporarily uneven transmission, reflecting differences in initial conditions across countries. Once the economic and financing conditions normalize, these transitory divergences dissipate and heterogeneity converges back to its baseline level. Heterogeneity then manifests itself as short-lived spikes.

The evidence presented in Figure 18 lends stronger support to the hypothesis of *cyclically driven heterogeneity*. First, the striking similarity in the evolution of dispersion across all variables suggests that the observed peaks are unlikely to reflect variable-specific factors operating differently across transmission channels. Rather, they point to a common underlying driver, most plausibly the monetary policy cycle itself. Second, and most importantly, the figure does not display the ratcheting-up pattern that would characterize structurally driven heterogeneity. Instead, periods of relatively low and stable dispersion alternate with sharp but transitory increases, and heterogeneity ultimately reverts toward a stable baseline level by the end of the sample.

Three episodes are particularly informative. The first sustained rise in heterogeneity coincides with the entry of March 2015 in the rolling window that marked the launch of APP in its full form⁵³ and the transition toward large-scale unconventional monetary accommodation. Subsequent rises in heterogeneity in our pattern align with the successive expansions of the programme via the different recalibrations that occurred following the programme's launch. Indeed, the use of the shadow rate enables us to capture the accommodation generated by unconventional instruments, notably of the APP since it captures the overall monetary policy stance beyond the zero lower bound. These asset purchases operate through multiple transmission channels, such as term premia compression and portfolio rebalancing, whose intensity varies across countries depending on their pre-existing conditions.

⁵³In March 2015, the ECB launched the Public Sector Purchase Programme (PSPP), the large-scale sovereign and public-sector bond purchases.

The programme's objective was to foster greater convergence in long-term interest rates across the monetary union and to ensure sufficiently accommodative long-term yields for all member states. This was achieved through disproportionately larger declines in borrowing costs for countries with initially elevated spreads,⁵⁴ thereby inducing asymmetric adjustments in other variables. Since our model incorporates the 10-year government bond spread, it explicitly captures this asymmetry: a common shadow rate shock translates into different paths across countries. Because heterogeneity is measured as the dispersion of these country-specific responses at given initial conditions, the APP's asymmetric action mechanically raises our metric, as illustrated in Figure 18, even though it is this asymmetric mechanism that has served the programme's objective of reducing cross-country divergence in monetary policy transmission. The second episode, characterized by an exceptionally low degree of heterogeneity during the second part of the ZLB phase, provides additional support for a cyclical interpretation. The third episode of interest emerges around the turning point of the 2022-2023 tightening cycle. Dispersion begins to rise in December 2021, coinciding with the ECB's shift toward a tightening stance amid rising inflation, before the first effective rate hike enters the rolling-window sample. A second marked increase then occurs after July 2022, the month of the initial hike, enters the rolling window. This pattern suggests that the speed and magnitude of the policy reversal over this period triggered a sharp but short-lived divergence in monetary transmission across countries. In each case, increases in dispersion are followed by a gradual normalization, suggesting that heterogeneity declines as the monetary policy cycle matures rather than becoming permanently entrenched.

Taken together, the figure documents temporary, policy-driven divergences in monetary transmission associated with periods of strong and/or unconventional intervention, followed by reversion toward a relatively low and stable baseline. The recurrence of pronounced spikes and subsequent normalization suggests that convergence in transmission resumes after active episodes, while heterogeneity remains compressed during calmer phases of the policy cycle. Such pattern is difficult to reconcile with the *structurally driven heterogeneity hypothesis*, which would rather imply a persistent upward drift compared to the baseline level of dispersion.

5.3.2 Share of variance explained by the *common component* across countries and over time

We now evaluate, at the country level, the importance of the common euro area component in driving the fluctuations of the variables included in the panel dataset. We first examine it over the full sample

⁵⁴This interpretation is further supported by the observation of the reaction of the 10-year sovereign bond yields following the launch of the APP measure in March 2015. Indeed, the launch of this unconventional monetary policy measure was associated with markedly heterogeneous responses from government bond yields across countries. In particular, they declined more in Portugal, Italy and Spain than in Germany (always displaying the lowest reaction). This observation is robust to alternative definitions of the pre- and post-announcement windows used to measure the reaction of the 10-year sovereign yields.

(Table 3), before turning to a rolling-window analysis to study how the contribution of the common component to the fluctuations of our data has evolved across countries and over time (Figure 19).

As implied by the measurement equation 5.1.2, each variable in the panel dataset X_t can be decomposed into a common component capturing, by definition, euro area-wide fluctuations and an idiosyncratic component specific to the series. The relative importance of these two components is allowed to vary across variables and over time, as each series loads differently on the vector of common factor summarizing euro area dynamics, both across variables and rolling windows for a given variable.

To quantify the contribution of the common component, we proceed as follows. For each posterior draw, we compute the R^2 from the regressions of each country-specific variable in X_t on the five estimated factors and the policy rate. This R^2 measures the share of variation in a given series explained by the common euro area component. Next, for each country, we average these R^2 values across variables to obtain a country-level measure of the explained variation. We do that at each draw. Finally, from the resulting distribution of R^2 across countries, we report the median draw along with its credible interval in Table 3 and Figure 19 below. Although this exercise does not directly address the question of the evolution of heterogeneity in monetary policy transmission *per se*, given that it does not involve any estimation of shock effects, it provides complementary insights into cross-country heterogeneity within the euro area.


First, it offers a diagnosis of our FAVAR framework by verifying that the identified common factors explain a substantial share of variation in individual series, and that their contribution is particularly important for euro area variables. Table 3 confirms, as expected, that over the full sample, it is the euro area common component that accounts for the largest share of variance in the euro area aggregate variables. Among individual countries, its contribution is highest in Spain, Italy, Finland, closely followed by France and Germany, whereas it is more limited in Belgium and Luxembourg. At the same time, Figure 19 also shows that over the whole period Luxembourg and the Netherlands exhibit the lowest contributions of the common component to the fluctuations of their variables, whereas Spain and the euro area aggregate (as expected) always lie at the upper end of the distribution. Despite the cross-country differences, the R^2 values remain within a narrow range over the whole sample, fluctuating between 0.65 and 0.90. Figure 19 also illustrates the increasing convergence over time for countries such as Portugal and Germany, suggesting a strengthening role of common euro area dynamics in shaping their macroeconomic, financial and banking fluctuations.

Table 3: Share of variance explained by the common component

Country	R^2	90% CI
EA	0.835	[0.833, 0.837]
ES	0.823	[0.816, 0.829]
IT	0.763	[0.759, 0.767]
FI	0.754	[0.750, 0.757]
FR	0.732	[0.729, 0.735]
NL	0.713	[0.709, 0.718]
PT	0.710	[0.707, 0.714]
AT	0.694	[0.691, 0.697]
DE	0.692	[0.688, 0.695]
BE	0.678	[0.675, 0.681]
LU	0.627	[0.624, 0.629]

Notes: The table reports the share of variation in the euro area and each country's variables explained by the common component (the five factors and the policy rate) over the full sample. At each posterior draw, the R^2 of the regression of each country-specific variable on the common component is averaged across country's variables to yield a single country-level R^2 ; the table reports the median of this measure across the retained draws, with the 90% credible interval (5th–95th percentiles). Higher values indicate a stronger role of common euro area dynamics in driving national fluctuations. Rows are ordered by decreasing R^2 .

Figure 19: Evolution of the share of variance explained by the common component



Notes: The figure plots, over a rolling window, the share of variance of euro area and each country's variables explained by the common component (the five estimated factors and the policy rate). At each posterior draw, we regress each country-specific variable on the common component and average the resulting R^2 across the country's variables to obtain a single country-level R^2 ; the figure reports the median R^2 across the retained draws, with shaded bands depicting the 68% and 90% credible intervals. The horizontal axis indicates the start and end dates of each rolling window; higher values denote a stronger role of the common euro area dynamics in driving national fluctuations.

In addition, this analysis can help shed light on which of the two hypotheses (*structurally versus cyclically driven heterogeneity*) better characterizes the observed heterogeneity in monetary policy transmission by exploring whether episodes of heightened heterogeneity are associated with a declining role of

the common euro area component (in favor of a *structurally driven heterogeneity*), or whether these two are largely uncorrelated (in favor of a *cyclically driven heterogeneity*). First, we observe that a direct one-to-one comparison between the evolution of monetary policy transmission heterogeneity documented in Figure 18 and of the share of the common component reported in Figure 19 is not straightforward. For each country, the evolution of the R^2 appears more stable than that of the degree of heterogeneity in monetary policy transmission. Moreover, the two sets of figures do not exhibit a clear mirror-image relationship: increases in monetary policy transmission heterogeneity are not systematically associated with or preceded by a decline in the share of variance explained by the common component. This absence of inverse co-movement lends support to the interpretation that the rise in transmission heterogeneity is not structurally driven. This observation holds for most countries for the first part of the chart, corresponding to the initial rise in monetary policy transmission heterogeneity, but matches the second peak less, around the post-COVID19 inflation surge and the most recent monetary tightening cycle, where the increase in monetary policy transmission dispersion in Figure 18 appears to coincide with a short-lived and slightly preceding decline in the share of variance explained by the common component for most countries.

5.3.3 Robustness to alternative specifications

Finally, we assess the sensitivity of our results to two alternative specifications. In our baseline specification, oil prices are included in the set of variables used to extract the common factors and classified as slow-moving. This assumption is consistent with treating the euro area as a small open economy with respect to global oil price dynamics, under which monetary policy shocks are assumed to exert no contemporaneous effect on oil prices. Under this specification, euro area monetary policy can influence oil prices only with a lag and primarily through the demand channel.⁵⁵ As a first robustness test, we assess whether our results are sensitive to an alternative specification in which oil prices are introduced directly into the VAR as an observed variable, ordered first, ahead of the common factors and the policy interest rate. Under this recursive identification scheme, the factors are allowed to respond contemporaneously to oil price shocks within the period. The policy rate, in turn, reacts within the same month to shocks originating from the factors as well as to oil price shocks. Conversely, oil price shocks are assumed not to respond contemporaneously to innovations in the policy rate. In the spirit of Boivin et al. (2009), this specification allows for the identification of an oil price shock alongside the monetary policy shock. The results of this alternative specification enable us to assess whether explicitly disentangling these two shocks affects our findings, given that oil price movements may independently influence inflation, output and financial conditions. Not distinguishing between them could result in external supply shocks being erroneously attributed to monetary policy innovations.

⁵⁵Oil prices are expressed in dollars, thereby abstracting from the exchange rate channel as a transmission mechanism following a euro area monetary policy shock.

However, the results of this specification, presented in the Appendix J.2, indicate that this distinction is not decisive since both the IRFs and the pattern of the evolution of heterogeneity are highly comparable to our baseline.

Second, we employ an alternative identification scheme for the monetary policy shock, drawing directly on high-frequency monetary policy surprises, in alignment with our local projection approach. To this end, we follow the high-frequency Cholesky VAR approach of Jarociński and Karadi (2020), which preserves the recursive identification structure while ordering the exogenous monetary policy surprise ahead of the policy rate, so that the monetary policy shock is identified through the response of the policy rate R_t to the surprise itself. This method lies midway between a Cholesky identification and a proxy VAR identification. Since our monetary policy surprise series is computed from the surprises in the 1-month, 3-month, 6-month and 1-year OIS rates, we chose to use the 1-year OIS rate as a policy rate proxy. The latter appears to be the most responsive interest rate to the monetary policy surprises. As in our baseline, financial variables are still classified as *fast moving* with respect to the policy rate. The results of this approach, reported in Appendix J.3, yield comparable and statistically significant impulse response functions, albeit a one-standard-deviation monetary policy shock in the series is associated with more muted macroeconomic effects compared to a shock of similar magnitude identified through Cholesky decomposition only. The resulting pattern of heterogeneity is consistent with a contained degree of heterogeneity in monetary policy transmission over time, ruling out any persistent increase thereof. Differences relative to our baseline results are discussed in Appendix J.3.

6 Conclusion

This paper reassesses the extent, the sources and the evolution of heterogeneity in monetary policy transmission (MPT) across euro area countries, twenty-five years after the monetary unification. It does so by proposing three complementary empirical strategies to offer a multidimensional assessment of heterogeneity in MPT. Overall, our findings suggest that the aggregate implications of cross-country divergences for MPT in the euro area should not be overstated.

First, we find, both visually and quantitatively, that the responses of HICP and real GDP exhibit a particularly low cross-country heterogeneity. Turning more specifically to the banking transmission channel of monetary policy, we observe that it operates in a broadly harmonized way across states. However, although prices of credit variables react rather similarly and remain tightly aligned with the euro area average, quantity-based credit measures display greater dispersion across countries, suggesting that, despite a rather homogeneous reaction of credit's prices, country-specific factors keep influencing the reactions of the demand and supply of loans granted by monetary and financial

institutions. Still, the responses of variables related to non-financial corporations are more homogeneous than those for households, reflecting deeper integration in corporate credit markets and the continued fragmentation of household mortgage markets that remain driven by national frameworks.

Second, our results show that household indebtedness, debt maturity, interest rate rigidity and sectoral composition significantly influence national responses to monetary policy shocks. However, their combined macroeconomic impact appears more nuanced. Indeed, we find that the dampening effect on MPT of lower exposure to interest rate fluctuations, due to a higher share of long-term maturity and/or of fixed-rate loans, can be offset by the amplifying effect of greater exposure to interest rate arising either via a higher share of flexible-rate loans, in one case and via a higher degree of household indebtedness, in the other. This same reasoning can be extended to other characteristics. Thus, our findings point to the existence of offsetting mechanisms among structural characteristics that may cancel each other's effects, helping reconcile the extensive evidence of country-level heterogeneity in MPT from the literature with the rather uniform aggregate responses of output and inflation. A key contribution of this paper is therefore to adopt a macro-level perspective that emphasizes the interaction of multiple structural features rather than isolating structural characteristics.

Third, we examine the time variation in MPT heterogeneity using a FAVAR framework. Overall, our measure of dispersion based on the full-sample estimation first confirms the particularly low cross-country heterogeneity in the transmission of monetary policy shocks to HICP and real GDP. Asymmetries do exist, but they lie elsewhere, in credit quantities, retail trade turnover, sovereign spreads, food prices. In addition, our rolling-window analysis reveals pronounced but episodic increases in heterogeneity of MPT over time that consistently revert to a low baseline level. These surges appear to coincide with periods of intensified policy intervention and to be largely driven by specific monetary policy cycles themselves rather than by structural factors reflecting deeper divergences. Our pattern notably illustrates how unconventional monetary policy tools, designed to address cross-country divergences, operate precisely by affecting euro area countries unevenly in order to restore that smoothness. This interpretation is supported by the observed desynchronization between peaks in MPT heterogeneity and the declines in the share of the variance of variables explained by the common component. Such evidence suggests that the observed heterogeneity does not stem from persistent and destabilizing structural divergences among member states.

Thus, the results presented in this paper convey a reassuring message. Despite existing structural differences across countries and the fact that the euro area does not fully satisfy the criteria of an OCA, MPT has remained smoother and more homogeneous across countries and over time than suggested by the existing literature.

References

- ALDER, M., N. COIMBRA, AND U. SZCZERBOWICZ (2023): “Corporate debt structure and heterogenous monetary policy transmission,” *Banque de France Working Paper*.
- ALLAYIOTI, A., L. GÓRNICKA, S. HOLTON, AND C. MARTÍNEZ HERNÁNDEZ (2024): “Monetary policy pass-through to consumer prices: evidence from granular price data,” *ECB Working Paper*.
- ALTAVILLA, C., L. BRUGNOLINI, R. S. GÜRKAYNAK, R. MOTTO, AND G. RAGUSA (2019): “Measuring euro area monetary policy,” *Journal of Monetary Economics*, 108, 162–179.
- ALTAVILLA, C., R. S. GÜRKAYNAK, AND R. QUAEDVLIEG (2024): “Macro and micro of external finance premium and monetary policy transmission,” *ECB Working Paper*.
- ANGELONI, A. AND M. EHRMANN (2003): “Monetary policy transmission in the euro area: any changes after EMU?” *ECB Working Paper*, 240.
- ANYFANTAKI, S., R. BARAHONA, D. CUCIC, D. FRICKE, P. HARTMANN, C. KAUFMANN, E. LUKMANOVA, AND A. MADDALONI (2026): “Monetary Policy Transmission and Non-Bank Financial Intermediation,” *ECB Occasional Paper*.
- ARTIS, M. J. (2003): “Reflections on the Optimal Currency Area (OCA) Criteria in the Light of EMU,” *International Journal of Finance and Economics*, 8(4), 297–307.
- BAI, J. AND S. NG (2002): “Determining the number of factors in approximate factor models,” *Econometrica*, 70(1), 191–221.
- BARIGOZZI, M., A. M. CONTI, AND M. LUCIANI (2014): “Do euro area countries respond asymmetrically to the common monetary policy?” *Oxford Bulletin of Economics and Statistics*, 76, 16–47.
- BAUER, M. AND E. T. SWANSON (2022): “A reassessment of monetary policy surprises and high-frequency identification,” *NBER Working Paper Series*, 29939.
- BAUMEISTER, C., P. LIU, AND H. MUMTAZ (2013): “Changes in the effects of monetary policy on disaggregate price dynamics,” *Journal of Economic Dynamics and Control*, 37(3), 543–560.
- BAYOUMI, T. AND B. EICHENGREEN (1992): “Shocking Aspects of European Monetary Unification,” *CEPR Discussion Papers*, 643.
- BERNANKE, B. S., J. BOIVIN, AND P. ELIASZ (2005): “Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach,” *Quarterly Journal of Economics*.
- BOIVIN, J., M. GIANNONI, AND B. MOJON (2009): “How Has the Euro Changed the Monetary Transmission Mechanism?” *NBER Macroeconomics Annual 2008*, 23, 77–125.
- BÉNASSY-QUÉRÉ, A. (2025): “Heterogeneity in the Euro Area and the Transmission of Monetary Policy,” *ECB Sintra Forum on Central Banking*.

- CALZA, A., T. MONACELLI, AND L. STRACCA (2013): "Housing Finance And Monetary Policy," *Journal of the European Economic Association*, 11, 101–122.
- CARLINO, G. AND R. DEFINA (1998): "The differential regional effects of monetary policy," *The Review of Economics and Statistics*, 80(4), 572–587.
- CARTER, C. K. AND R. KOHN (1994): "On Gibbs Sampling for State Space Models," *Biometrika*, 81, 541–553.
- CHITU, L., S. GORI, AND R. S. GÜRKAYNAK (2025): "External Finance Premium: Market Finance versus Bank Finance," *ECB Working Paper*.
- CHOW, G. C. AND A.-L. LIN (1971): "Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series," *The Review of Economics and Statistics*, 53(4), 372–75.
- CHRISTOFFEL, K., K. KUESTER, AND T. LINZERT (2009): "The role of labor markets for euro area monetary policy," *Working Paper Series*, 1035.
- CIAPANNA, E. AND C. RONDINELLI (2014): "Retail Market Structure and Consumer Prices in the Euro Area," *ECB Working Paper Series*.
- CICCARELLI, M., A. MADDALONI, AND J.-L. PEYDRÓ (2013): "Heterogeneous transmission mechanism: monetary policy and financial fragility in the eurozone," *Economic Policy*, 28, 459–512.
- CORSETTI, G. (2008): "A modern reconsideration of the theory of Optimal Currency Area," *European Economy - Economic Papers*, 308.
- CORSETTI, G., J. DUARTE, AND S. MANN (2022): "One Money, Many Markets" (Monetary Transmission and Housing Financing in the Euro Area)," *Journal of the European Economic Association*, 20, 513–548.
- CUMMING, F. AND P. HUBERT (2022): "House prices, the distribution of household debt and the refinancing channel of monetary policy," *Economics Letters*, 212, 110280.
- DATSENKO, R. AND J. FLECK (2024): "Country-Specific Effects of Euro-Area Monetary Policy: The Role of Sectoral Differences," *Fed Notes*.
- EHRMANN, M. (2000): "Comparing Monetary Policy Transmission across European Countries," *Weltwirtschaftliches Archiv*, 56–83.
- FALCK, E., M. HOFFMANN, AND P. HÜRTGEN (2021): "Disagreement about inflation expectations and monetary policy transmission," *Journal of Monetary Economics*, 118, 15–31.
- FLODÉN, M., M. KILSTRÖM, J. SIGURDSSON, AND R. VESTMAN (2017): "Household Debt and Monetary Policy: Revealing the Cash-Flow Channel," *Sveriges Riksbank Working Paper Series*.
- GEORGIADIS, G. (2015): "Examining asymmetries in the transmission of monetary policy in the euro area: Evidence from a mixed cross-section global VAR model," *European Economic Review*.


- GOLINSKI, A. AND P. SPENCER (2024): “Unconventional Monetary Policies and the Yield Curve: Estimating Non-Affine Term Structure Models with Unspanned Macro Risk by Factor Extraction,” *The Review of Asset Pricing Studies*, 14(1), 119–152.
- GONÇALVES, S., A. M. HERRERA, L. KILIAN, AND E. PESAVENTO (2024): “State-dependent local projections,” *Journal of Econometrics*, 244 (2).
- HAUPTMEIER, S., F. HOLM-HADULLA, AND E. SAPORITO (2025): “Economic Structure and the Real Effects of Monetary Policy,” *Applied Economics Letters*.
- HEDLUND, A., K. LARKIN, K. MITMAN, AND S. OZKAN (2025): “Mortgage Market Structure and the Transmission of Monetary Policy During the Great Inflation,” *Fed Notes*.
- HERBST, E. P. AND B. JOHANNSEN (2024): “Bias in local projections,” *Journal of Econometrics*, 240.
- HOLM-HADULLA, F. AND S. POOL (2025): “Interest Rate Control and the Transmission of Monetary Policy,” *ECB Working Paper Series*.
- HOLM-HADULLA, F. AND C. THÜRWÄCHTER (2021): “Heterogeneity in corporate debt structures and the transmission of monetary policy,” *European Economic Review*, 136, 103743.
- INOUE, A., O. JORDÀ, AND G. M. KUERSTEINER (2024): “Inference for Local Projections,” *Federal Reserve Bank of San Francisco Working Paper*, 29.
- JAROCIŃSKI, M. AND P. KARADI (2020): “Deconstructing Monetary Policy Surprises – The Role of Information Shocks,” *American Economic Journal: Macroeconomics*, 12, 1–43.
- JORDÀ, O. (2005): “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- (2023): “Local Projections for Applied Economics,” *Federal Reserve Bank of San Francisco Working Paper*.
- JORDÀ, O. AND A. TAYLOR (2025): “Local projections,” *Journal of Economic Literature*, 63 (1), 59–110.
- KOROBILIS, D. (2012): “Assessing the Transmission of Monetary Policy Using Time-varying Parameter Dynamic Factor Models,” *Oxford Bulletin of Economics and Statistics*, 75(2), 157–179.
- LENZA, M. AND G. E. PRIMICERI (2022): “How to estimate a vector autoregression after March 2020?” *Journal of Applied Econometrics*, 37(4), 688–699.
- LUETTICKE, R. (2021): “Transmission of Monetary Policy with Heterogeneity in Household Portfolios,” *American Economic Journal: Macroeconomics*, 13, 1–25.
- MANDLER, M., M. SCHARNAGL, AND U. VOLZ (2022): “Heterogeneity in Euro Area Monetary Policy Transmission: Results from a Large Multicountry BVAR Model,” *Journal of Money, Credit and Banking*, 54, 627–649.

- MIRANDA-AGRIPPINO, S. AND G. RICCO (2021): “The Transmission of Monetary Policy Shocks,” *American Economic Journal: Macroeconomics*, 13, 74–107.
- MONTIEL OLEA, J. L., M. PLAGBORG-MØLLER, E. QIAN, AND C. K. WOLF (2025): “Local Projections or VARs? A Primer for Macroeconomists,” *Working Paper*.
- MUMTAZ, H. (2010): “Evolving UK Macroeconomic Dynamics: A Time-Varying Factor Augmented VAR,” *Bank of England Working Paper*, 386.
- MUNDELL, R. (1961): “A Theory of Optimum Currency Areas,” *American Economic Review*, 51(4), 657–665.
- OUERK, S., C. BOUCHER, AND C. LUBOCHINSKY (2020): “Unconventional monetary policy in the Euro Area: Shadow rate and light effects,” *Journal of Macroeconomics*, 65.
- PEERSMAN, G. AND F. SMETS (2002): “The industry effects of monetary policy in the euro area,” *Working Paper Series*, 165.
- PICA, S. (2022): “Housing Markets and the Heterogeneous Effects of Monetary Policy Across the Euro Area,” *Working Paper*.
- PIGER, J. AND T. STOCKWELL (2023): “Differences from Differencing: Should Local Projections with Observed Shocks be Estimated in Levels or Differences?” *SSRN Working Paper*.
- RAMEY, V. AND S. ZUBAIRY (2018): “Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data,” *Journal of Political Economy*, 126, 850–901.
- SCIACOVELLI, G. (2025): “Monetary Policy Transmission Through Adjustable-Rate Mortgages in the Euro Area,” *Working Paper*.
- SLACALEK, J., O. TRISTANI, AND G. L. VIOLANTE (2020): “Household Balance Sheet Channels of Monetary Policy: A Back of the Envelope Calculation for the Euro Area,” *NBER Working Paper*.
- WONG, A. (2021): “Refinancing and the transmission of monetary policy to consumption,” *American Economic Review*.

A Construction of the euro area monetary policy shocks

Below, we compare our monetary policy surprise series—constructed by replicating and extending the surprises of [Jarociński and Karadi \(2020\)](#)—with the authors’ original series (MPI median), which is obtained by purging the surprises identified by [Altavilla et al. \(2019\)](#) from the *central bank information* effect using the sign-restriction VAR approach described in their paper. The two series are very close and strongly overlap. The two main discrepancies are concentrated around two specific dates, which in our case reflect the inclusion of market reactions to the exceptional ECB interventions of 11/09/2001 and 08/10/2008.

Figure A.1: Monetary policy instrument



B Local projections data description


Below, we present the data sources and transformations of our local projection approach, variable by variable. The detailed econometric specifications of the local projections are available upon request.

Table B.1: Data sources and transformations for the local projections

Variable	Description	Transformation
Real GDP	Interpolated with the Chow–Lin decomposition method at a monthly frequency, using the Industrial Production Index from the Short-Term Statistics dataset and the quarterly real GDP from the Main Aggregates in the National Accounts dataset.	$(\log(GDP_t) - \log(GDP_{t-1})) \times 100$
HICP	Monthly HICP Overall Index from the Indices of Consumer Prices dataset. The series have been seasonally adjusted using the X-13ARIMA-SEATS Seasonal Adjustment Program from Matlab.	$(\log(HICP_t) - \log(HICP_{t-1})) \times 100$
Unemployment rate	Monthly unemployment rate for individuals aged 15 years and over as part of the Infra-Annual Labour Force Statistics (IALFS) from OECD.	$UNRATE_t - UNRATE_{t-1}$
NFC borrowing cost	Monthly cost of borrowing of new loans granted to non-financial corporations by deposit-taking institutions except the central bank. The data comes from the MFI Interest Rate Statistics Dataset.	R_t^{firms}
HH borrowing cost for housing purchase	Monthly cost of borrowing of new loans granted to households and non-profit institutions for house purchase by deposit-taking institutions except the central bank. The data comes from the MFI Interest Rate Statistics Dataset.	$R_t^{mortgages}$
Growth rate of loans to firms	Monthly annual growth rate of loans to firms built using the outstanding stock of loans held at period-end by MFIs excluding the ESCB and granted to non-financial corporations. The data comes from the Balance Sheet Items Dataset.	$\left(\frac{Loans_t^{firms} - Loans_{t-12}^{firms}}{Loans_{t-12}^{firms}} \right) \times 100$
Growth rate of loans to HH	Monthly annual growth rate of loans to HH built using the outstanding stock of loans held at period-end by MFIs excluding the ESCB and granted to HH. The data comes from the Balance Sheet Items Dataset.	$\left(\frac{Loans_t^{HH} - Loans_{t-12}^{HH}}{Loans_{t-12}^{HH}} \right) \times 100$

C Comparison of EA responses with Driscoll-Kraay standard errors

Figure C.1: Comparison of the EA responses to a 1-std monetary policy shock



Notes: The figure compares different euro area impulse responses to a one-standard-deviation monetary policy shock across four specifications. EA (aggregate data) uses area-wide series; EA (panel – 11 countries) relies on a balanced panel of the eleven founding members; EA (panel – 19 countries) uses all current member states except Croatia and Bulgaria; and EA (GDP-weighted – all countries) is the GDP-weighted average of the country-specific responses, each weighted by its share of euro area GDP. Shaded areas depict one-standard-deviation confidence bands based on DK standard errors. No bands are reported for the EA (GDP-weighted – all countries), essentially shown for informational purposes, since the corresponding estimates are derived ex post by averaging the country-specific IRFs.

D Alternative local projection estimation: country-specific responses recovered from a panel estimation

D.1 Impulse responses from the panel local projection estimation

In Figure D.1 below are presented the country-specific IRFs to a one-standard deviation monetary policy shock recovered from the panel fixed effect estimation (3.1.3). We have decomposed the euro


area average response β_h^{EA} obtained from the usual panel specification into country-specific effects by interacting the monetary policy shock with country dummies. As a reminder, $\beta_{c_{ref}}$ captures the country-specific response of the reference country (Austria, in alphabetical order), while for each other country c at horizon h , the response to a monetary policy shock is given in deviation from the reference country. Thus, $\beta^{c,h} = \beta_{c_{ref}}^{c,h} + \delta^{c,h}$.

$$y_{t+h} = \alpha^h + \sum_{c \neq c_{ref}} \gamma_c D_c + \beta_{c_{ref}}^{c,h} \cdot MPS_t + \sum_{c \neq c_{ref}} \delta^{c,h} (MPS_t \times D_c) + \sum_{k=1}^p \phi_k^{h,c} \cdot X_{t-k}^c + \epsilon_{i,t+h} \quad (D.1.1)$$


By construction, computing the unweighted mean of these country-specific effects yields the same value as β_h^{EA} directly estimated in the fixed-effects panel specification (3.1.2). Although the country-specific impulse responses recovered from the panel specification share some common patterns to those obtained from the country-by-country local projections in specification (3.1.1), they are not identical. Their comparison is limited as these two series of national estimates are based on a different information set. Indeed, this panel specification exploits the cross-country correlations, completely absent from the country-by-country set of regressions obtained with specification (3.1.1). Whereas the country-by-country local projections estimate the effect of the euro-area monetary policy shock on each country in isolation of the others and ignores cross-country co-movements and spillovers effects, the panel local projection jointly uses all cross-sectional variation. This can naturally lead to certain differences in the estimated transmission mechanism. Since all countries are exposed to the same euro-area monetary policy shock, it can be justified to treat them as part of a unified system responding to a common disturbance. By better capturing the idea of one shock and multiple responses this specification therefore brings additional information, on top of providing material for robustness checks in the computation of the CV.

Figure D.1: Country-specific panel local projections IRFs to a 1-std monetary policy shock

(a) Transmission to macroeconomic variables



(b) Transmission to the banking variables




Notes: The figure plots the country-specific IRFs to a one-standard-deviation monetary policy shock, recovered from the panel fixed-effects specification (3.1.3) that interacts the shock with country dummies to decompose the euro area average effect into country-specific effects. Each country response is obtained as $\beta^{c,h} = \beta^{c,ref,h} + \delta^{c,h}$; the EA panel data is computed as the unweighted average across $\beta^{c,h}$. Shaded areas represent 68% confidence bands around the EA aggregate response, computed using clustered standard errors (dark blue) and the conservative Driscoll-Kraay standard errors (light blue).

D.2 Testing the equality of country-specific estimates from panel local projections

In this section, we report the results of pairwise tests of equality between country-specific impulse responses one year after a monetary policy shock. The effects are obtained from the panel local projection presented above. The test results are consistent across specifications using clustered and Driscoll–Kraay standard errors: in both cases, we do not reject the null hypothesis of equality for any pair of countries and any variables after accounting for multiple testing correction. We report the results based on Driscoll–Kraay standard errors, which correspond to the most conservative specification.

Figure D.2: Pairwise tests of equality of country-responses estimated from panel LP





Notes: Each cell reports the t-statistic of the pairwise test of equality between the horizon-12 responses of two countries, $H_0 : \beta_i(12) = \beta_j(12)$. Responses are estimated with the country-by-country local projections from specification (3.1.1). Lighter cells corresponds to smaller t-statistics; a value below 1.96 indicates that the two responses cannot be distinguished at the 5% significance level in an individual test. Red dots would flag pairs for which equality is rejected at the 5% family-wise level, after correcting for multiple testing via the Holm (1979) procedure; no such rejection occurs for any variable, which is why no red dot appears.

E Explaining the sources of heterogeneities



E.1 Country characteristics driving heterogeneity in MPT

The data sources and variables used in this section are described in Section 4.1.



Figure E.1: Country structural characteristics tested as potential drivers of national responses to a monetary policy shock



Share of loans over gross disposal income




Ratio of manufacturing-to-services value added



E.2 Testing significance with Driscoll-Kraay standard errors

In this section, we report our results using Driscoll-Kraay standard errors, which are robust to cross-sectional dependence, serial correlation and heteroskedasticity, following [Inoue et al. \(2024\)](#). Despite wider confidence intervals, we find that the statistical significance of most of our economic effects still holds.


Figure E.2.1: The role of fixed versus variable interest rate loans to firms and households




Notes: The left charts report, the effect of a one-std monetary policy shock when the characteristics is at its sample mean; the middle charts report, the marginal effect associated with a 1 std increase in the characteristic above its mean. The right charts display the total effect of the shock, evaluated at the 25th and 75th percentiles of the cross-country distribution of the characteristic, where each country's value is averaged over the sample period. This averaging is conservative, as it smooths out within-country variation over time and therefore narrows the gap between the two evaluation points.

Figure E.2.2: The role of the maturity of loans to firms and households

(a) Specification 1




(b) Specification 2




Notes: The left charts report, the effect of a one-std monetary policy shock when the characteristics is at its sample mean; the middle charts report, the marginal effect associated with a 1 std increase in the characteristic above its mean. The right charts display the total effect of the shock, evaluated at the 25th and 75th percentiles of the cross-country distribution of the characteristic, where each country's value is averaged over the sample period. This averaging is conservative, as it smooths out within-country variation over time and therefore narrows the gap between the two evaluation points.

Figure E.2.3: The role of the level of household indebtedness

(a) Specification 1




(b) Specification 2




Notes: The left charts report, the effect of a one-std monetary policy shock when the characteristics is at its sample mean; the middle charts report, the marginal effect associated with a 1 std increase in the characteristic above its mean. The right charts display the total effect of the shock, evaluated at the 25th and 75th percentiles of the cross-country distribution of the characteristic, where each country's value is averaged over the sample period. This averaging is conservative, as it smooths out within-country variation over time and therefore narrows the gap between the two evaluation points.

Figure E.2.4: The role of the sectoral composition

(a) Specification 1




(b) Specification 2




Notes: The left charts report, the effect of a one-std monetary policy shock when the characteristics is at its sample mean; the middle charts report, the marginal effect associated with a 1 std increase in the characteristic above its mean. The right charts display the total effect of the shock, evaluated at the 25th and 75th percentiles of the cross-country distribution of the characteristic, where each country's value is averaged over the sample period. This averaging is conservative, as it smooths out within-country variation over time and therefore narrows the gap between the two evaluation points.

Figure E.2.5: Evidence of offsetting effects across characteristics using Driscoll-Kraay


(a) Interaction between debt maturity and interest rate regime



(b) Interaction between debt maturity and the level of HH indebtedness



(c) Interaction between interest rate regime and the level of HH indebtedness



Notes: Each panel reports results from specification (4.3.1), in which the monetary policy shock is interacted simultaneously with two structural characteristics, $\tilde{Z}_1$ and $\tilde{Z}_2$, both standardized and lagged by one period. Within each panel, the left plot shows the effect of a one-standard-deviation monetary policy shock when both characteristics are at their euro area averages; the middle plot shows the marginal effect of the first characteristic; and the right plot shows the triple-interaction coefficient, capturing the additional effect on transmission when both characteristics are simultaneously one standard deviation above their averages. Vertical bars depict the one standard deviation confidence interval built using Driscoll-Kraay standard errors.

E.3 Assessing the magnitude of the offsetting mechanism

When both $\tilde{Z}_{1,t-1}$ and $\tilde{Z}_{2,t-1}$ are one standard deviation above their respective means ($\tilde{Z}_{1,t-1} = 0$ and $\tilde{Z}_{2,t-1} = 0$), the total effect of the monetary policy shock becomes $\beta_h^1 + \beta_h^4 + \beta_h^5 + \beta_h^6$. So the offsetting between the two characteristics of interest is total if $\beta_h^4 + \beta_h^5 + \beta_h^6 = 0$. The tables below report the corresponding p-values of the Wald test, horizon by horizon, for the three cases tested, with clustered as well as Driscoll-Kraay standard errors:

- **Interaction between LT debt maturity and variable rate loans (Tables E.3.1 and E.3.2 and Tables E.3.3 and E.3.4)**

With clustered standard errors:

Table E.3.1: HICP

H	p-value	H	p-value	H	p-value
H0	0.90	H7	0.21	H14	0.20
H1	0.71	H8	0.31	H15	0.16
H2	0.52	H9	0.27	H16	0.17
H3	0.74	H10	0.34	H17	0.13
H4	0.43	H11	0.44	H18	0.12
H5	0.38	H12	0.33		
H6	0.30	H13	0.16		

Table E.3.2: Real GDP

H	p-value	H	p-value	H	p-value
H0	0.57	H7	0.26	H14	0.05
H1	0.35	H8	0.08	H15	0.12
H2	0.44	H9	0.11	H16	0.63
H3	0.29	H10	0.06	H17	0.36
H4	0.26	H11	0.31	H18	0.44
H5	0.67	H12	0.31		
H6	0.98	H13	0.15		

With Driscoll-Kraay standard errors:

Table E.3.3: HICP

H	p-value	H	p-value	H	p-value
H0	0.92	H7	0.08	H14	0.09
H1	0.76	H8	0.16	H15	0.07
H2	0.48	H9	0.15	H16	0.08
H3	0.68	H10	0.20	H17	0.07
H4	0.36	H11	0.33	H18	0.07
H5	0.25	H12	0.17		
H6	0.15	H13	0.06		

Table E.3.4: Real GDP

H	p-value	H	p-value	H	p-value
H0	0.75	H7	0.52	H14	0.18
H1	0.40	H8	0.53	H15	0.20
H2	0.48	H9	0.42	H16	0.77
H3	0.48	H10	0.33	H17	0.54
H4	0.38	H11	0.45	H18	0.59
H5	0.81	H12	0.35		
H6	0.99	H13	0.25		

- **Interaction between LT debt maturity and high degree of HH indebtedness (Tables E.3.5 and E.3.6 and Tables E.3.7 and E.3.8)**

With clustered standard errors:

Table E.3.5: HICP

H	p-value	H	p-value	H	p-value
H0	0.13	H7	0.05	H14	0.14
H1	0.03	H8	0.06	H15	0.19
H2	0.05	H9	0.04	H16	0.28
H3	0.01	H10	0.05	H17	0.37
H4	0.02	H11	0.09	H18	0.33
H5	0.07	H12	0.09		
H6	0.07	H13	0.18		

Table E.3.6: Real GDP

H	p-value	H	p-value	H	p-value
H0	0.78	H7	0.49	H14	0.60
H1	0.55	H8	0.41	H15	0.79
H2	0.49	H9	0.04	H16	0.69
H3	0.04	H10	0.03	H17	0.89
H4	0.02	H11	0.39	H18	0.54
H5	0.72	H12	0.35		
H6	0.15	H13	0.91		

With Driscoll-Kraay standard errors:

Table E.3.7: HICP

H	p-value	H	p-value	H	p-value
H0	0.48	H7	0.12	H14	0.31
H1	0.22	H8	0.12	H15	0.43
H2	0.22	H9	0.10	H16	0.51
H3	0.06	H10	0.15	H17	0.57
H4	0.07	H11	0.23	H18	0.57
H5	0.16	H12	0.21		
H6	0.14	H13	0.34		

Table E.3.8: Real GDP

H	p-value	H	p-value	H	p-value
H0	0.71	H7	0.60	H14	0.69
H1	0.64	H8	0.57	H15	0.88
H2	0.59	H9	0.13	H16	0.85
H3	0.22	H10	0.40	H17	0.97
H4	0.24	H11	0.61	H18	0.87
H5	0.81	H12	0.51		
H6	0.25	H13	0.93		

- Interaction between fixed-rate loans and high degree of HH indebtedness (Tables E.3.9 and E.3.10 and Tables E.3.11 and E.3.12)

With clustered standard errors:

Table E.3.9: HICP

H	p-value	H	p-value	H	p-value
H0	0.24	H7	0.03	H14	0.10
H1	0.03	H8	0.06	H15	0.08
H2	0.00	H9	0.02	H16	0.06
H3	0.00	H10	0.02	H17	0.07
H4	0.01	H11	0.06	H18	0.08
H5	0.02	H12	0.04		
H6	0.03	H13	0.18		

Table E.3.10: Real GDP

H	p-value	H	p-value	H	p-value
H0	0.50	H7	0.38	H14	0.75
H1	0.30	H8	0.36	H15	0.90
H2	0.28	H9	0.37	H16	0.65
H3	0.09	H10	0.79	H17	0.81
H4	0.45	H11	0.64	H18	0.90
H5	0.40	H12	0.26		
H6	0.13	H13	0.74		

With Driscoll-Kraay standard errors:

Table E.3.11: HICP

H	p-value	H	p-value	H	p-value
H0	0.47	H7	0.20	H14	0.17
H1	0.32	H8	0.30	H15	0.15
H2	0.18	H9	0.09	H16	0.14
H3	0.10	H10	0.08	H17	0.14
H4	0.11	H11	0.13	H18	0.15
H5	0.08	H12	0.10		
H6	0.13	H13	0.27		

Table E.3.12: Real GDP

H	p-value	H	p-value	H	p-value
H0	0.65	H7	0.23	H14	0.78
H1	0.67	H8	0.25	H15	0.90
H2	0.61	H9	0.30	H16	0.62
H3	0.14	H10	0.79	H17	0.82
H4	0.40	H11	0.69	H18	0.92
H5	0.28	H12	0.34		
H6	0.04	H13	0.79		

F FAVAR Data description

Table F.1 summarizes the data sources, the series transformations and the specific treatments applied in estimating our FAVAR model. The transformation applied to each series is indicated by Index 1, where 1 denotes the year-on-year growth rate, 2 the level (no transformation) and 3 the 12-month difference. The classification of variables as fast- or slow-moving is reported under Index 2, where 0 indicates a slow-moving variable and 1 a fast-moving variable. All data are seasonally adjusted.

Table F.1: Data description and transformation

Series	Source	Index 1	Index 2
Real GDP	Eurostat, Main aggregates, National Accounts	1	0
Unemployment rate	OECD, Infra-Annual Labour Force Statistics	3	0
Retail trade turnover	Eurostat, Short Term Statistics	1	0
HICP all items	Eurostat, Indices of Consumer Prices	1	0
HICP energy	Eurostat, Indices of Consumer Prices	1	0
HICP food	Eurostat, Indices of Consumer Prices	1	0
Producer Price Index, domestic sales	Eurostat, Short Term Statistics	1	0
10-year gvt bond spread	OECD, Long-Term Government Bond Yields: 10-Year	2	1
Stock market prices	OECD, Financial Market: Share Prices	1	1
Households credit	ECB, Balance Sheet Items	1	0
Firms credit	ECB, Balance Sheet Items	1	0
Oil prices	Bloomberg L.P., Brent Crude (ICE), CO1 Comdty	1	0
Nominal effective exchange rate-17/Euro	ECB, EXR	1	1
EURIBOR 3M (shadow rate)	Banque de France, internal computations	2	–

G Factors selection

G.1 Factor loadings


Table G.1: The top ten series loading the most heavily on each factor.

Factor 1	Factor 2	Factor 3	Factor 4	Factor 5
Producer prices NL	Spread FI	non-MFI credit PT	Retail trade ES	Stock prices ES
Producer prices FR	Spread AT	HH credit FI	HICP energy ES	Stock prices BE
Producer prices EA	Spread NL	HH credit PT	Unemployment IT	Stock prices EA
HICP energy AT	Stock prices AT	non-MFI credit FI	Unemployment ES	Stock prices DE
HICP energy EA	Retail trade turnover EA	non-MFI credit ES	Real GDP ES	Stock prices FR
Producer prices PT	HICP food FI	HH credit ES	Unemployment EA	Stock prices AT
Producer prices DE	HICP all items FI	non-MFI credit IT	HICP food NL	Stock prices AT
Producer prices FI	Spread LU	non-MFI credit EA	HICP all items FR	Stock prices PT
HICP all items LU	Retail trade turnover FI	Stock prices NL	Spread PT	Stock prices IT
Producer prices AT	Spread BE	non-MFI credit NL	Real GDP EA	Stock prices LU

G.2 Factor sensitivity analysis

This subsection investigates the sensitivity of our IRFs to the number of factors included in the FAVAR. We find that there exists a sensitivity across the number of factors. This sensitivity is driven by the inclusion of the 10-year government bond spreads among the variables. In a more parsimonious specification that excludes all country-specific spread variables, the full-sample IRFs are fully comparable whether four, five or six factors are used and the estimates are fully stable across factor choices again, as presented in Figure G.1 below. For conciseness, we report here, as an example, only the factor sensitivity analysis results for the euro area aggregate variables included in the panel. Corresponding country-level results are available upon request.


Figure G.1: Factor sensitivity analysis





H Country-by-country IRFs with confidence intervals


Figure H.1 presents our baseline IRFs for each variable and each country individually, with their 68% and 90% credible intervals. We also report the IRFs for the global variables and for the policy rate.


Figure H.1: Country-by-country IRFs to a 1 standard deviation monetary policy shock











I FAVAR-based comparison of heterogeneity in MPT across samples

Table I.1 compares the coefficient of variation obtained on our baseline sample (2000–2025) with that obtained on the sample used by Corsetti et al. (2022) (2000–2016), holding our FAVAR specification, variables, iterations, draw selection and identification fixed, at the exception of the number of factors since our selection criteria retain six to seven factors on this shorter sample. Heterogeneity in MPT is higher on the shorter sample for most variables (value in bold), while the ranking across variables is broadly preserved. This suggests that the sample period is a driver of the lower heterogeneity documented in our baseline results. The magnitude of this reduction in heterogeneity when adding the 2016–2025 period varies considerably across variables. For real GDP, the shift is modest, in line with existing literature already pointing to a lower degree of heterogeneity for responses of real GDP. In contrast, responses of credit and price variables display a substantial decline in heterogeneity.

Table I.1: Comparison of the degree of heterogeneity in MPT across subsamples

Variable	CV (2000–2025)	CV (2000–2016)*
Real GDP	0.399	0.42
Unemployment rate	0.748	0.63
Retail trade turnover	1.800	7.35
HICP all items	0.476	1.95
HICP energy	0.336	0.57
HICP food	2.243	2.05
PPI	0.282	0.63
Spread	1.051	2.86
Stock prices	0.194	0.33
Households credit	1.221	6.47
non-MFI credit	1.034	1.52

* The CVs for the 2000–2016 subsample are computed using $K = 6$ factors, the minimum number of factors to use on this shorter sample. The reported values and the difference relative to the full-sample CVs are fully robust to using 7 factors.


J Rolling window exercise: robustness checks

J.1 Sensitivity to window length

We replicate the exercise presented in Section 5.3.1, substituting the 12-year rolling window employed in the baseline specification with an 11-year window.⁵⁶ As illustrated in Figure J1.1, the principal

⁵⁶The 11- and 12-year windows are reported on the chart as 12- and 13-year windows (2000–2012 and 2000–2013 respectively), since year-on-year growth rates require one extra year of level data to compute the first observation.

Figure J1.1: Time variation in cross-country dispersion of IRFs (11-years window length)




Notes: Each panel plots, for a given variable, the evolution over time of cross-country heterogeneity in the transmission of a monetary policy shock, estimated from a rolling window FAVAR model. For each window, heterogeneity is measured as the cross-country standard deviation of the impulse responses, averaged across horizons. The two lines allow a comparison across rolling-window lengths: one uses a 12-year window (black, baseline), the other a shorter 11-year window (blue). The horizontal axis reports the start and end dates of each rolling window; shaded bands depict the 68% and 90% credible intervals. Vertical dashed lines mark key monetary policy events.


pattern characterizing the evolution of monetary policy transmission heterogeneity proves fully robust to this alternative window length. It also documents a generally subdued degree of heterogeneity, punctuated by transitory increases coinciding with key phases of the monetary policy cycle — notably periods of unconventional monetary policies or strong monetary policy interventions. However, the blue heterogeneity series displays a higher volatility, manifested through more erratic peaks or discontinuities. This pattern is the consequence of the reduced sample size, rather than a genuine shift in the underlying dynamics of heterogeneity. Shorter estimation windows inherently introduce greater noise and translate, for certain sub-windows, into convergence failures at early draws of the Gibbs sampling algorithm, thereby preventing the computation of IRFs and the subsequent quantification of cross-sectional heterogeneity. These considerations motivated the adoption of a longer estimation window in the baseline specification.

J.2 Sensitivity to an alternative FAVAR specification with an oil price shock

The results presented in Figures J2.1 and J2.2 correspond to a specification in which oil prices are directly included in the VAR as an observed exogenous variable and ordered first, ahead of the common factors and the policy interest rate. This specification enables the identification of an additional structural shock—an oil price shock—alongside the monetary policy shock. The heterogeneity pattern under this new specification remains highly consistent with the baseline results.


Figure J2.1: IRFs to a one-standard-deviation monetary policy shock, full-sample estimation






Notes: The figure reports the country-specific IRFs to a one-standard-deviation monetary policy shock, recovered from the FAVAR specification where oil prices are included as an observed variable ordered first, ahead of the common factors and the policy rate. The model is estimated at monthly frequency, on a balanced panel of 121 series covering the euro area and ten founding members over January 2000–July 2025, excluding the COVID-19 period. Responses are expressed in percentage points. For each country and variable, the solid lines report the median response across draws.

Figure J2.2: Time variation in cross-country dispersion of IRFs






Notes: Each panel plots, for a given variable, the evolution over time of cross-country heterogeneity in the transmission of a monetary policy shock, estimated from a FAVAR model over rolling windows. For each window, heterogeneity is measured as the cross-country standard deviation of the impulse responses, averaged across horizons. The two lines allow a comparison across model specifications: one includes an additional oil price shock (blue), the other is our baseline (black). The horizontal axis reports the start and end dates of each rolling window; shaded bands depict the 68% and 90% credible intervals. Vertical dashed lines mark key monetary policy events.

J.3 Sensitivity to an alternative monetary policy shock identification

This section presents the results of FAVAR estimation after identifying the monetary policy shock from high-frequency monetary policy surprises following [Jarociński and Karadi \(2020\)](#). The pattern of heterogeneity obtained under this alternative identification scheme conveys the same overall message as our baseline specification: heterogeneity remains low over time, with no evidence of persistent divergence in the effects of monetary policy shocks across euro area countries that would be the symptom of a structural fragmentation. During the post-covid inflation surge and the subsequent monetary tightening cycle, we observe a slightly moderate increase in heterogeneity, particularly for price-related variables, while GDP responses remain particularly stable. However, the identification based on monetary policy surprises produces a flatter heterogeneity profile for all variables. Two factors may account for this. First, the surprise-based identification uses the one-year OIS rate as the policy rate, since the OIS rates respond more strongly than the shadow rate to the high-frequency surprises, which are themselves constructed from OIS rates. No shadow rate is used as policy rate. As a result, the heterogeneity peaks observed in the baseline and associated with episodes of non-standard monetary policy are not fully captured with this alternative specification. Second, monetary policy surprises are quantitatively small and exhibit limited volatility over a substantial part of the sample, reflecting the high degree of anticipation of ECB policy decisions. Consequently, shocks identified


via the effect of high-frequency surprises on the one-year OIS rate induce comparatively modest movements in the policy rate and generate both weaker macroeconomic responses and dispersion. Indeed, a one-standard-deviation shock constructed from monetary policy surprises implies a much smaller change in the policy rate compared to a one-standard-deviation shock identified under the baseline recursive identification scheme.

Figure J3.1: IRFs to a 1-std monetary policy shock identified with monetary policy surprises, full sample



Notes: The figure reports the country-specific IRFs to a one-standard-deviation monetary policy shock identified with high-frequency monetary surprises, following [Jarociński and Karadi \(2020\)](#). The model is estimated at monthly frequency, on a balanced panel of 121 series covering the euro area and ten founding members over January 2000–July 2025, excluding the COVID-19 period. Responses are expressed in percentage points. For each country and variable, the solid lines report the median response across draws.

Figure J3.2: Time variation in cross-country dispersion of IRFs



Notes: Each panel plots the evolution of cross-country heterogeneity in the transmission of a monetary policy shock, estimated from a FAVAR model over rolling windows. The monetary policy shock is identified with high-frequency monetary policy surprises. For each window, heterogeneity is measured as the cross-country standard deviation of the impulse responses, averaged across horizons. The horizontal axis reports the start and end dates of each rolling window; shaded bands depict the 68% and 90% credible intervals.