



Federal Reserve
Bank of Dallas

The Short Lags of Monetary Policy

Gergely Buda, Vasco M. Carvalho, Giancarlo Corsetti, João B. Duarte,
Stephen Hansen, Afonso S. Moura, Álvaro Ortiz, Tomasa Rodrigo,
José V. Rodríguez Mora and Guilherme Alves da Silva

Working Paper 2631

September 2026

Research Department

<https://doi.org/10.24149/wp2631>

Working papers from the Federal Reserve Bank of Dallas are preliminary drafts circulated for professional comment. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

The Short Lags of Monetary Policy*

Gergely Buda[†], Vasco M. Carvalho[‡], Giancarlo Corsetti[§], João B. Duarte[‡],
Stephen Hansen[¶], Afonso S. Moura[°], Álvaro Ortiz[□], Tomasa Rodrigo^{*},
José V. Rodríguez Mora⁺ and Guilherme Alves da Silva¹

September 2, 2026

Abstract

We examine the transmission of monetary policy shocks to the macroeconomy at high frequency. To do so, we construct daily consumption and investment aggregates from bank transaction records and leverage administrative data on daily corporate sales and employment for Spain. We show that variables typically viewed as slow-moving, such as consumption and corporate sales, respond significantly within weeks. By contrast, aggregate employment and (monthly) consumer prices adjust only gradually. Disaggregating responses by sector, consumption category, and supply-chain distance from final demand, we find that fast demand adjustment is led by downstream sectors closely tied to consumption—especially luxury goods and durables—while upstream sectors respond more slowly but more persistently, ultimately accounting for the contraction in aggregate activity. Excess inventories accumulate first in retail and wholesale trade and only later among upstream manufacturers, reconciling the rapid adjustment of sales with the slower response of production and employment. Finally, we show that time aggregation to the quarterly frequency obscures the timing of monetary policy transmission by shifting estimated responses to longer horizons, whereas weekly and monthly aggregation preserve the patterns observed at daily frequency.

Keywords: Event-study, Monetary Policy, Economic Activity, High-Frequency Data, Local Projections

JEL Classification: E31, E43, E44, E52, E58

* The authors thank, without implicating, George-Marios Angeletos, Adrien Auclert, Michael Bauer, Ambrogio Cesa-Bianchi, Daniel Dias, Thomas Drechsel, Fabrizia Mealli, Stefano Pica, Giorgio Primiceri, Ricardo Reis, Antonio Rua, Guido Spano, Vincent Sterk, Oreste Tristani, Mark Watson, Ivan Werning, Christian Wolf and seminar participants at Bank of Belgium, Bank of England, Bank of Portugal 11th Conference on Monetary Economics, CEBRA Annual Meeting, CEPR-ESSIM, ECB-EIB-Bank of Italy Conference on the Use of Surveys for Monetary Policy, Federal Reserve Board, Federal Reserve Bank of San Francisco, Hong Kong University, London Business School, Società Italiana degli Economisti, University of Cambridge, University College London for helpful suggestions and discussions. We thank Juan Ramón García for providing the daily employment data and Sirenia Vazquez for her work on the investment series. Carvalho gratefully acknowledges funding from ERC Consolidator Grant 101001221. Duarte and Moura gratefully acknowledge funding from Fundação para a Ciência e a Tecnologia (2022.07354.PTDC, UID/00124/2025, UID/PRR/124/2025, PTDC/EGE-ECO/7620/2020, and CEECIND/03227/2018), and LISBOA2030 (DataLab2030-LISBOA2030-FEDER-01314200). Hansen gratefully acknowledges financial support from ERC Consolidator Grant 864863. An early version of this paper was circulated with the title “Short and Variable Lags”. The views expressed in this paper do not necessarily reflect the views of the Eurosystem, the Federal Reserve Bank of Dallas or the Federal Reserve System.

[†]Gergely Buda, Barcelona School of Economics. [‡]Corresponding author: Vasco M. Carvalho, University of Cambridge and CEPR, ympmdc2@cam.ac.uk. [§]Giancarlo Corsetti, European University Institute and CEPR. [‡]João B. Duarte, Nova School of Business and Economics. [¶]Stephen Hansen, UCL, FRB Dallas, and CEPR. [°]Afonso S. Moura, European Central Bank and Banco de Portugal. [□]Álvaro Ortiz, BBVA Research and CRIW. ^{*}Tomasa Rodrigo, BBVA Research. ⁺José V. Rodríguez Mora, CUNEF Universidad, U. of Edinburgh and CEPR. ¹Guilherme Alves da Silva, Nova School of Business and Economics.

"Monetary actions affect economic conditions only after a lag that is both long and variable" (Friedman, 1961).

1 Introduction

Milton Friedman's dictum is, to this day, firmly ingrained in the thinking of both academics and policymakers. Decades of research and policy experience have shown that the transmission mechanism of monetary policy is no doubt complex, playing out over multiple channels and fully unfolding at medium to long horizons, as Friedman emphasized early on. Complementing Friedman's dictum, a widely held view in policy circles holds that monetary policy transmission can be envisioned as a two-stage process: at first, it affects asset prices, financial conditions and expectations; over time, these feed through to key real aggregates and inflation with increasing intensity. This view challenges theoretical monetary models, where in standard specifications the largest response of both financial and real variables to monetary policy takes place on impact—motivating a variety of frictions (such as adjustment costs) and/or behavioral elements (such as habit formation) that help rationalize delays in the real effects of policy measures (see the discussion in, e.g., [Woodford 2003](#) and [Mackowiak and Wiederholt 2009](#)).

In this paper, we assemble a novel, comprehensive daily dataset for Spain and use it to reassess the conventional wisdom on the lags of monetary policy transmission. By combining these data with monetary policy shocks identified from high-frequency financial market surprises, we estimate local projections that align the frequency of the shocks with that of the outcome variables, thereby avoiding the biases introduced by temporal aggregation. Specifically, we study the response of household expenditure, corporate sales to both households and firms, investment, and employment in the days, weeks, and months following a monetary policy shock, at both the aggregate and disaggregated levels, over a one-year horizon. Our granular analysis reveals that: (i) the impact of monetary policy shocks on aggregates that are typically considered "slow-moving", such as aggregate corporate sales, consumption and investment, can be detected within days and weeks, rather than quarters or years; (ii) this fast adjustment is driven by a contraction in final demand, and the sales response of downstream sectors closest to it, particularly by producers of durable and luxury consumption goods, whereas the response by upstream sectors is slower, but overall more persistent; (iii) the fast adjustment of demand detectable at daily, weekly and monthly frequencies is no longer detectable when using quarterly frequency data (common in macroeconomic studies), because of time-aggregation bias; (iv) the timing of adjustment of financial and real variables points to a well-defined transmission sequence—information about the shock

reaches households almost immediately through interest-rate news and the repricing of credit conditions; together with a deterioration in expectations, this triggers a rapid contraction in demand, most pronounced among mortgagors but not confined to them; the contraction in final demand then propagates gradually to production, with inventories buffering the gap between fast-adjusting sales and slower-adjusting output as demand moves from downstream to upstream sectors; and (v) the short lags in the response of demand translate into a slow, smooth but progressively deeper decline in employment and prices—consistent with the idea that monetary policy affects these key variables with long lags. Our results therefore suggest that these long lags primarily reflect the gradual transmission of an initially rapid contraction in demand, rather than a slow response of demand itself to monetary policy shocks.

Articulating our main results in greater detail, our first contribution is the construction and deployment of novel high-quality daily proxies for four key macroeconomic aggregates in Spain: aggregate corporate sales, aggregate consumption, aggregate investment and aggregate employment. To this end, we leverage unfettered access to both retail and corporate transactions recorded by one of the largest private banks in Spain and combine it with newly available tax and social security data tabulations by the Spanish government. In particular, we construct daily consumption and investment series from the bottom up, from bank transactions associated with the universe of, respectively, household and corporate accounts of Banco Bilbao Vizcaya Argentaria (BBVA). We complement this with newly assembled daily data on corporate sales to both households and firms—which provide a proxy for daily gross output up to changes in inventory stocks—and aggregate employment as recently compiled by, respectively, the Spanish Tax Authority based on VAT declarations, and the Spanish Ministry for Inclusion, Social Security, and Migration in the course of tracking public pension obligations. The breadth and granularity of the underlying data also allow us to construct a rich set of disaggregated series, which we complement with traditional high-frequency financial indicators and a broad range of monthly macroeconomic variables.

Based on our daily series, our second contribution consists of documenting that monetary policy has economically and statistically significant effects on real economic activity already within weeks from a policy innovation. Major components of demand and corporate sales, conventionally classified as “slow moving” (Bernanke, Boivin and Elias, 2005), closely track the fast response of financial variables and expectations. Our baseline results show that firm sales decline steadily from the first weeks after a contractionary monetary policy shock—in the unsmoothed daily estimates, the contraction is statistically detectable within the first week—reaching a local trough of -0.50% 51 days

after a one standard deviation shock and stabilizing thereafter, before declining again roughly 200 days after the monetary policy shock. Consumption follows a broadly similar pattern, with an initial trough at 44 days after the monetary policy shock, followed by a spell of stabilization and then again a contraction at longer lags. Relative to sales, the consumption response is smaller, with an initial trough at -0.39% . Investment responds along similar lines—relative to these variables, however, the investment response is noisier yet still detectable at short lags, and more persistent and stronger at longer horizons, around eight months after the shock. In turn, the aggregate employment response, while statistically detectable early on, is initially very contained and, relative to the other three variables, is smoother and steadier. Its strongest response within the first year is precisely at the end of our analysis’ horizon, on day 365, with a cumulative decline of -0.22% . These patterns highlight a sharp difference in lags: final demand reaches economically meaningful troughs within the first two months of the shock, whereas employment builds toward its largest response only at the one-year horizon.

Turning to disaggregated series, our third contribution consists of documenting how monetary shocks propagate across categories of consumption, corporate sales by sector and by sector upstreamness. We show that the response of categories of goods traditionally considered more responsive to monetary policy—such as durables and luxuries, e.g., transport, clothing, and furnishing—is not only sharp, but also fast: it is significant at very short lags. The same pattern characterizes services such as health, education, and restaurants. Conversely, we find no short lags in the response of housing services and utilities, communication, and food and beverages (“food at home”), which remain insignificant, or are even positive, across all horizons. Building on [Antràs et al. \(2012\)](#), we classify sectoral activity depending on the position of the sector in the production network, distinguishing between upstream and downstream sectors. We show that the response of downstream sectors, more closely tied to final demand, is both faster (statistically significant in about one month) and much larger (three times larger at peak) when compared to upstream sectors. Also mirroring fast-moving final demand, downstream sales stabilize in the second quarter, before contracting further at long lags. In contrast, the contraction of upstream sectors, providing general purpose inputs for the production of goods and services, is somewhat slower, becoming statistically significant 60 days after the shock, but much more persistent. The differences in these responses suggest a pattern of gradual upstream transmission of an initial downstream, final demand adjustment to a monetary policy impulse—lending empirical support to an emerging view that production networks and supply chains may be critical in the transmission of monetary policy.

Our fourth contribution is to provide empirical evidence on the “time aggregation bias”, the subject of a long-standing foundational literature analyzing whether and why the infrequent measurement of economic variables may distort empirical inference. Exploiting the unique features of our high-frequency dataset, we show that aggregating our daily data into quarterly frequency—the frequency at which macro variables have commonly been available to researchers—alters the empirical responses of sales, consumption and investment to monetary policy shocks, obscuring economically meaningful dynamics. At the same time, we find that aggregating daily into weekly and monthly frequencies preserves qualitative features of daily-frequency results. In other words, a researcher with access to only a quarterly-frequency aggregation of our data would not be able to detect a within-quarter adjustment—which would be instead detectable with data at the daily, weekly or monthly frequency. Said researcher would therefore wrongly conclude that consumption and corporate sales only start reacting “slowly” in the second quarter following the shock. This finding suggests that sharp classifications of variables as slow and fast moving in response to monetary shocks may be misdirected by the frequency of data most commonly available in macroeconometric work. A Monte Carlo exercise, calibrated to our estimated daily responses and to the actual calendar of ECB announcements, reproduces these patterns in a controlled environment where the true response is known. The simulation establishes that the distortion is not an artifact of our short quarterly sample but originates in the time aggregation of the monetary policy shocks within the quarter.

Our fifth contribution brings this evidence to bear on a consistent narrative of the transmission, extending the analysis with newly constructed daily series and a large set of monthly variables. We start by presenting direct evidence on how quickly information about monetary policy reaches households. We construct a novel daily indicator of interest rate news across hundreds of Spanish media outlets and show that coverage sharply increases and its tone becomes negative within days of a monetary policy shock. Over the same, very short horizon, indicators of credit conditions and house prices deteriorate; and expectations and sentiment surveys fall in tandem, so that the observed actions taken by agents match both the timing and content of their information updates. We then exploit customer-level variation in mortgage status—merging BBVA’s credit register with our transaction-level consumption data—and document that mortgagors, directly exposed to rising borrowing costs, cut consumption by more than non-mortgagors.¹

Next, assembling monthly data on inventories from Spain, we document that inven-

¹Nonetheless, we also report a significant, rapid response of non-mortgagors, suggesting the relevance of a complementary, expectations channel: consumption demand contracts fast, regardless of the household position in the credit market.

tories adjust to absorb the gap that opens as final sales contract ahead of production, in line with their long-documented role as buffers (Blinder and Maccini, 1991; Ramey and West, 1999; Bils and Kahn, 2000). Mirroring the sectoral propagation patterns we observe in sales, downstream retail firms accumulate inventories first, as consumer demand falls on impact while shipments from manufacturers continue; manufacturers accumulate stocks of finished products later, once retailers cut their orders.² The timing is internally consistent with the response of employment, which begins falling substantially when manufacturers report excess inventories and reduce production.

We conclude our analysis by focusing on the two key variables of interest to central banks. We show that monthly CPI inflation—both in the aggregate and across expenditure categories—responds only gradually to monetary policy shocks, much like monthly employment. We further show that the sluggish adjustment of aggregate employment largely reflects the dynamics of workers on permanent contracts. Thus, our evidence indicates that prices and permanent-contract employment are the last variables to adjust, responding with lags that are consistent with both the existing empirical literature and the conventional policy view of monetary transmission.

Overall, these results call into question conventional views, holding that there is a sharp distinction between slow-moving real variables and fast-moving financial variables, and envisioning two sequential, financial and real, stages in the transmission of monetary policy, as discussed by, e.g., Burr and Willems (2024). Our high-frequency evidence reveals a different timing and ordering of the adjustment that refines rather than overturns Friedman’s dictum: monetary policy does affect employment and prices only with long and variable lags, but those lags arise because a rapid contraction in demand propagates only gradually through production, rather than because demand itself responds slowly to the policy impulse.

Results from extensive robustness exercises are discussed in the final section of the paper and the appendix. In particular, in light of our results on time aggregation—suggesting that monthly frequency analysis largely preserves the response patterns obtained with our daily series—we use our rich monthly dataset to assess the response of alternative data sources and proxies of output, investment and consumption in longer samples. We also demonstrate robustness to a wide variety of alternative seasonality and smoothing procedures, as well as to alternative monetary policy shock series. Finally, we deploy alternatives to our simple local projections baseline methodology and consider robustness to alternative treatments of the COVID-19 pandemic period.

²Specifically, inventory growth in wholesale and retail trade is positive on impact and turns down only after about three months, while survey measures of excess stocks rise first among retailers (in the first month), and only later among manufacturers (around the third month).

Literature. Our paper relates to several strands of the literature. The first is the rapidly growing empirical literature that studies the transmission of monetary policy using shocks identified at high frequency around policy announcements, an approach pioneered by Kuttner (2001) and Gürkaynak, Sack and Swanson (2005). In our paper, we use results in Jarociński and Karadi (2020), who construct a database of monetary policy surprises for the Euro Area. Our key novel contribution is to align daily-frequency monetary shocks with daily-frequency measures of economic activity.³ We estimate empirical responses avoiding time-aggregation of either real variables or the shocks to lower frequencies, thereby avoiding time-aggregation bias.

The second strand is the recent literature that exploits naturally occurring transaction-level data to track economic dynamics at high frequency (see, e.g., Andersen et al. (2021), Andersen et al. (2022), Bounie et al. (2020), Buda et al. (2022), Chetty et al. (2020), and Ganong and Noel (2019)).⁴ Among studies of household consumption, closest to ours is Grigoli and Sandri (2026), who provide evidence on the rapid and sizable credit card expenditures’ responses to monetary policy shocks for Germany. Importantly, this study suggests that our evidence of short transmission lags is not sample or country-specific.⁵ These authors compare the effects of short- and medium-term monetary surprises, as well as rate hikes and cuts, and trace the effects by subcategories of consumption goods and proxies for household income. Our study differs in two crucial respects. First, our consumption measure is constructed from the universe of bank transactions, rather than from a single payment instrument, and closely tracks the consumption aggregate in the Spanish national accounts, following the validation exercise in Buda et al. (2022).⁶ Second, our evidence extends well beyond household consumption: as described above, our dataset covers firms’ sales, investment and employment decisions as well as household expenditures, which allows us to trace the transmission of monetary policy across

³The literature has routinely aggregated monetary surprises to lower frequencies, including monthly (Gertler and Karadi (2015); Almgren et al. (2022)), quarterly (Cloyne, Ferreira and Surico (2020)), and even annual frequencies (Holm, Paul and Tischbirek (2021)).

⁴These contributions are in turn related to a rapidly expanding literature focused on high-frequency indicators of economic activity, motivated—especially during the COVID-19 pandemic—by the need to support policy decisions in a rapidly changing environment. Examples of weekly indicators include Eraslan and Götz (2021), Baumeister, Leiva-León and Sims (2021), and Lewis et al. (2022), while Diebold (2022) and Rua and Lourenço (2020) develop daily measures of economic activity.

⁵Concerning short lags, we should also note that recent studies using monthly data for the euro area and the United States have found same-month response in industrial production and other (often interpolated) measures of economic activity, see Jarociński and Karadi (2020) and Miranda-Agrippino and Ricco (2021).

⁶Specifically, our high-frequency consumption measure is constructed following the methodology in Buda et al. (2022). When aggregated to quarterly frequency, it closely matches the household consumption series in the Spanish national accounts. Another study using the same data is Ferreira et al. (2022), which examines the effects of inflation on households’ balance sheets.

households, firms, and sectors within a unified empirical framework.

Crucially, our paper relates to a smaller but foundational literature on the consequences of time aggregation. Early seminal contributions on the theoretical properties of econometric models with temporally aggregated data include [Amemiya and Wu \(1972\)](#), [Sims \(1971\)](#), and [Geweke \(1978\)](#), while [Marcet \(1991\)](#) studies the implications of temporal aggregation for forecasting. The relevance of time-aggregation bias for a classic empirical question in macroeconomics—whether money growth Granger-causes inflation—is examined by [Christiano and Eichenbaum \(1987\)](#) and [Stock \(1987\)](#). More recently, [Jacobson, Matthes and Walker \(2026\)](#) revisit this issue and show that temporal aggregation bias is an important quantitative source of the “price puzzle”—the rise in inflation following a monetary contraction often found in empirical studies—and provides a complementary explanation to mechanisms such as the Fed information channel. Consistent with this evidence, we find no evidence of a price puzzle when using monthly inflation data. Our distinct contribution is to offer systemic empirical evidence of significant time aggregation bias in estimating the lags of monetary policy across a number of key real outcomes.

Furthermore, our evidence on the mechanisms underlying monetary transmission connects to at least four additional strands of the literature. To start, our analysis builds on a consolidated body of studies showing that financial markets react essentially instantaneously to monetary policy shocks ([Bernanke and Kuttner, 2005](#); [Gürkaynak, Sack and Swanson, 2005](#); [Hanson and Stein, 2015](#); [Swanson, 2021](#)) and that interest rates directly relevant for households, such as mortgage rates, react within weeks ([Gorea, Kryvtsov and Kudlyak, 2023](#)). More recently, a growing literature has challenged the traditional view that households are largely inattentive to monetary policy developments: [Lewis, Makridis and Mertens \(2019\)](#) document that public confidence in the state of the economy responds *immediately* to surprises about the Federal Funds target rate; [Lamla and Vinogradov \(2019\)](#) show that central bank announcements significantly raise the probability that consumers receive monetary policy news within days; and [Coibion, Gorodnichenko and Weber \(2022\)](#) show that household inflation expectations respond immediately when policy communication reaches them. Our daily news indicators complement this literature by shedding light on the supply side of the information channel. We show that media coverage of interest rates rises sharply, while its tone becomes markedly negative, within days of an ECB monetary policy shock. This evidence is consistent with the epidemiological model of [Carroll \(2003\)](#), in which news media act as the conduit through which macroeconomic developments are transmitted to households.

The second strand is the literature on the mortgage cash-flow channel: [Cloyne, Fer-](#)

reira and Surico (2020) show that the household spending response to monetary policy is concentrated among mortgagors; Di Maggio et al. (2017) find that automatic resets of adjustable-rate mortgages pass through to household spending; Flodén et al. (2021) show that indebted households cut consumption significantly more than debt-free households following an increase in interest rates; and Calza, Monacelli and Stracca (2013) and Corsetti, Duarte and Mann (2022) relate the strength of the consumption response across countries to the structure of housing finance, highlighting the prevalence of variable-rate mortgages, common in Spain, as a key determinant of monetary transmission. Consistent with this channel, we find that mortgagors reduce consumption substantially more than non-mortgagors. At the same time, however, we document a significant and rapid response of non-mortgagors, suggesting that monetary policy also operates through channels beyond direct interest-rate exposure.

The third strand of literature concerns the role of inventories in business cycle fluctuations and the transmission of monetary policy (Blinder and Maccini, 1991; Ramey and West, 1999; Bils and Kahn, 2000; Bernanke and Gertler, 1995). Combining daily data on sales with monthly information on inventories and business surveys, we provide a granular account of their role and dynamics in the transmission mechanism across the production network. We show that in the wholesale and retail sectors inventories begin to accumulate immediately after a monetary policy shock, then start to decline after roughly three months—at that point, manufacturers first report excess stocks.

Finally, our findings speak to the literature on production networks and monetary transmission. Building on the seminal contribution of Basu (1995), a growing theoretical and quantitative literature—including Carvalho (2006), Nakamura and Steinsson (2010), and, more recently, Pasten, Schoenle and Weber (2020)—has developed multi-sector frameworks to study how heterogeneous price rigidities and input-output linkages shape the propagation of monetary policy shocks. Closer in spirit to our analysis, Ozdagli and Weber (2017) and Ghassibe (2021) provide empirical evidence on the role of production networks in amplifying monetary disturbances. Focusing on the response of asset prices, Ozdagli and Weber (2017) uncover that demand contractions originating in downstream sectors spread to upstream suppliers. Our contribution differs from this study in that we document propagation along the production network via the response of real economic activity across sectors, rather than the response of financial variables.

The remainder of the paper is organized as follows. Sections 2 and 3 describe the data and the empirical methodology. Section 4 presents our baseline findings, documenting the short-lags in the response of aggregate real activity to monetary policy shocks. Section 5 explores the heterogeneity of these responses across consumption and sales cate-

gories, and upstream versus downstream sectors. Section 6 discusses time aggregation. Section 7 presents extensions and robustness analyses using monthly data and discusses an empirically grounded account of the transmission mechanism. Section 8 concludes.

2 Data

We assemble what is, to the best of our knowledge, the first dataset with daily, high-quality measures of aggregate corporate sales, consumption, investment and employment. To do this, we first leverage access to the universe of retail and corporate transaction records of one of the largest banks in Spain, BBVA, and construct novel series of daily aggregate consumption, daily aggregate investment, as well as daily consumption demand disaggregated by COICOP categories. Second, we complement this unique data with administrative daily series recently compiled by Spanish ministries and tax authorities: daily aggregate employment from Spanish Social Security records, and both aggregate and sectoral daily sales from Value Added Tax declarations by firms to the Spanish Tax Authority. Third, we further enrich our data by compiling more standard daily series on interest rates and stock returns as well as new, daily indicators of interest rate news coverage in Spanish media. At the monthly level, we additionally compile data on consumer prices, housing prices, and forward-looking expectations, confidence indexes and inventories.⁷

2.1 Daily Economic Activity Data

In this section we detail the data sources and methods underlying the construction of our daily measures of real economic activity in Spain, which proxy for daily corporate sales, private consumption, private investment, and employment. We also review a host of monthly counterparts to our baseline daily measures. Although our baseline series have varying starting dates, they all end in October 2023. Summary statistics for the various daily series are provided in Appendix A.2.

2.1.1 Aggregate Sales

Data on aggregate corporate sales at daily frequency are publicly available through the Spanish Tax Authority. The Tax Authority compiles the series from daily Value Added Tax (VAT) declarations by firms, reporting their domestic sales transactions—the tax base

⁷Table A1 in the appendix provides an overview of measures, sources and time coverage for the data used in the paper.

for VAT—on each day.⁸ This variable encompasses sales of final consumption goods to Spanish households and tourists, as well as sales of investment goods to Spanish firms and domestic firm-to-firm intermediate transactions. It can thus be taken as a proxy for daily gross domestic output, up to changes in inventories.

Only large firms or conglomerates—specifically, those with a turnover of 6 million euros or above in the previous year—are legally required to supply their domestic sales information daily. According to the Spanish Tax Authority ([Agencia Tributaria, 2023](#)), the number of firms reporting daily sales in 2019 was approximately 60,000 (out of a universe of 3.8 million VAT-paying entities in Spain). Nevertheless, due to their size, these firms accounted for about 70% of domestic sales by all firms in the same year.

Based on the same information, the Spanish Tax Authority also releases daily sales series disaggregated by NACE sector. The available series account for at least 50% of each sector-level sales—with the exceptions of ‘Hospitality Services’ and the residual, catch-all, sector labelled ‘Remaining Activities’, as a larger fraction of economic activity in these sectors is accounted for by smaller firms that are not obliged to report daily to the tax authority. Because of their low representativeness, we drop these two sectors from our analysis; see Appendix [A.2](#) and [Agencia Tributaria \(2023\)](#) for details.

Following the recommendations of the Spanish Tax Authority, we deflate daily series using monthly price indexes. Specifically, we apply the same monthly deflator value to all observations in the month.⁹ We deflate aggregate sales using the Spanish Consumer Price Index (CPI); Manufacturing and Construction sector sales using the respective producer price indexes (PPI); Wholesale and Retail Trade, and Transportation and Storage with appropriate disaggregated CPIs. Finally, we use the Services Price Index (SPI) to deflate the remaining (service) sectors.

Note that, based on the same VAT reporting data, the Spanish Tax Authority additionally compiles a range of *monthly* series, which we also employ in our analysis. These monthly series have two advantages. First, while the earliest available date for the daily sales series is July 1st, 2017, the monthly series start much earlier, in January 2000. Second, at monthly frequency, the Spanish Tax Authority provides sectoral breakdowns not available at daily frequency. Most importantly, it offers a breakdown of domestic

⁸With the introduction of the VAT immediate declaration system (“Suministro Inmediato de Información”), large taxpayers included in the system are required to send the Tax Agency the details of the billing records within four days of the issuance of an invoice. While comprehensive for most of Spain, this reporting system excludes firms with activity exclusively within the province of Navarra or the Basque Country, who report to their own regional tax agencies. In addition, activities in territories with no VAT – the Canaries, Ceuta and Melilla – are not represented.

⁹As an alternative, we also experimented with linear interpolation of the monthly price index series—the resulting real daily series are indistinguishable from our baseline.

corporate sales by use – distinguishing monthly domestic consumption, investment and intermediate input sales – supplementing it with a monthly export series.^{10,11}

2.1.2 Consumption

We construct a proxy for daily aggregate consumption using the universe of bank transactions recorded in the Spanish retail accounts of Banco Bilbao Vizcaya Argentaria (BBVA)—following the methods developed by Buda et al. (2022). In particular, our measure is a daily counterpart to their proxy for quarterly and annual aggregate consumption of private households. Below we give a brief overview of the construction of our series and refer the reader to Buda et al. (2022) for further detail.

The underlying data source for our consumption data is the universe of bank account outflows—including all card transactions, cash withdrawals, regular direct debits and occasional transfers—of Spanish residents who hold a retail account with BBVA. These account outflow records are supplemented by extensive metadata on both bank clients and individual transactions. The baseline sample consists of transactions for 1.8 million BBVA ‘active customers’—defined as bank clients that made at least ten consumption-related transactions in each quarter of the sample—excluding individuals who are self-employed.

Based on this data, the construction of a proxy for aggregate consumption involves two main steps, as detailed by Buda et al. (2022), that we follow closely. First, not every account outflow of these customers corresponds to a consumption expenditure: Buda et al. (2022) classify individual transactions as consumption, savings, investment or tax payments relying on the metadata associated to them, and following the national accounting principles from the European System of Accounts. The exceptions are cash withdrawals, which are assumed to serve exclusively for consumption expenditures.¹² Further, and again following European System of Accounts’ recommendations, Buda et al. (2022) impute housing services to all customers. Second, since the population of BBVA retail customers differs from the Spanish adult population along observables, Buda et al. (2022) aggregate consumption by summing over individual consumption using Spanish population weights at the gender-age-neighborhood cell level.

¹⁰According to Agencia Tributaria (2023) (see the Table “Destination of Domestic Production at Basic Prices,” in the National Accounts), the Spanish Tax Authority classifies domestic sales by calculating the proportion of the demand for intermediates, final consumption expenditure, and gross capital formation, based on national accounts data and in particular, the input-output matrix.

¹¹In a further robustness check, we also use a standard monthly industrial production series for Spain as a proxy for manufacturing output.

¹²Analogously, difficulties in classifying the purpose of transactions by self-employed people explain why this group is excluded from our sample.

While we follow closely Buda et al. (2022), our focus on daily frequencies requires us to take a stand on regular consumption expenditures happening at lower frequencies than daily. This includes monthly (imputed) housing services and the payment of regular utility bills, often on regular direct debits. We distribute these regular expenditures uniformly across all days of the month, assuming a regular service flow to households.

Overall, as Buda et al. (2022) show, the aggregates implied by this large scale consumption panel match well the official quarterly aggregate consumption series in both levels and growth rates: the implied level of aggregate consumption is, on average, within 1% of its official national accounts counterpart and the correlation of quarter-on-quarter growth rates across the two series is 0.987.

Further, exploiting metadata associated with each consumption transaction, Buda et al. (2022) show how to construct category-specific consumption series following the European COICOP system and distinguishing 11 consumption categories, which we also use here; see Appendix A.2 for a brief description of these consumption categories. As shown in Buda et al. (2022), the implied distribution for category-specific consumption shares matches well their official national accounts' counterpart.

We also merge the transaction-level consumption data with BBVA's credit register, which covers the universe of mortgage contracts held at the bank and is updated monthly. For this purpose, rather than aggregating consumption across customers using sampling weights, as we do when constructing the aggregate daily consumption series, we retain the data at the customer-day level. We then sum each customer's daily expenditures within each calendar month to obtain total monthly consumption at the customer level and merge the resulting customer-month panel with the credit register using customer identifiers. The matched panel runs from July 2020 to May 2025 and comprises approximately 1.7 million customers and 72 million customer-month observations. The merged dataset allows us to distinguish, in each month, between customers who hold a mortgage with BBVA and those who do not.

As with daily sales series, consumption daily series are deflated using monthly price indexes—aggregate consumption is deflated with the CPI, while individual consumption category series are deflated using CPI at the COICOP level. Our sample starts in August 1st, 2015 and ends in October 30th, 2023. Finally, as mentioned above, note that we also have access to an alternative *monthly* consumption series, as assembled by the Spanish Tax Authority from VAT data; see the discussion in the previous subsection. This monthly data serves as an important robustness check, as it is compiled from wholly distinct and publicly available data source.

2.1.3 Investment

Our daily aggregate investment proxy is also constructed using transaction data from BBVA. Our starting point is now the universe of corporate accounts and corporate transactions at the bank. In particular, we have access to all corporate transfers, and all (reverse) factoring operations mediated by BBVA. The former corresponds to all direct firm-to-firm payments while the latter, also known as confirming, is a form of supply chain finance service provided by the bank.

From this universe, we extract transactions that can be reasonably inferred to be a payment for goods and services. We do this by restricting to transactions that include the word ‘invoice’—or variations thereof—in metadata associated with each transaction. Further, we only keep operations where we can identify both parties—sender and receiver—as firms. We implement this either because we can identify both parties as BBVA corporate customers or because we can identify the non-BBVA party in the Spanish Business Registry.¹³ We further drop transactions involving financial sector firms, public sector firms or entities listed as non-profit foundations and eliminate all transactions where both parties belong to the same ownership group. Finally, note that while we are able to date each confirming transaction according to the original corporate invoice date, transactions stemming from direct firm-to-firm transfers are dated according to the date of the transfer of funds; they may therefore lag the original contract between firms. The resulting dataset comprises 17.4 million transactions—roughly evenly split between corporate transfers and reverse factoring operations—among 1.9 million distinct corporate entities, occurring between April 2017 and October 2023.¹⁴

We are then faced with two hurdles. First, as is the case with retail customers, the population of BBVA corporate customers is a biased sample of the Spanish population of firms: we find that the sectoral coverage of BBVA does not coincide with the sectoral distribution of output in Spain. To address this problem, we apply sector-level weights such that the aggregate of yearly (total) corporate sales by sector within the BBVA sample matches the respective sectoral aggregate—yearly domestic sales of investment goods plus domestic sales of intermediate goods—as recorded in Spanish Input-Output tables, for each year.

The second hurdle is the difficulty in ascertaining whether a transaction between two firms corresponds to a trade in investment goods and services, as opposed to in-

¹³Information is available through the standard SABI – Iberian Balance Sheet Analysis System – firm-level dataset, which includes the near universe of all Spanish firms.

¹⁴Among the entities in our sample, about 1.1 million firms are BBVA corporate clients, while the remaining are externally matched via the corporate registry.

intermediate inputs. This is because, unlike the case of consumption above, the metadata associated to each transaction is not sufficient to classify corporate transactions according to use. To make progress, we broadly follow the solution proposed by the Spanish Tax Authority, which faces the same hurdle in compiling monthly investment series from corporate VAT data. Namely, we exploit Spanish national accounts' input-output tables to obtain, for each sector, the share of investment goods sales in total sales. We then apply this share to all BBVA corporate transactions involving all firms in a sector, as sellers of goods and services, on any given day. Summing this across all firms and sectors gives our proxy for daily aggregate investment. Appendix A.3 presents the steps involved in constructing this daily investment series in more detail.

We benchmark our aggregate investment proxy by time-aggregating it at the monthly and quarterly frequency and comparing the resulting series with two series released by the Spanish authorities: quarterly national accounts investment and the above described monthly aggregate investment series published by the Spanish Tax Authority from VAT declarations, which we also use in the paper. As discussed in detail in Appendix A.3 our investment series tracks these two alternative lower frequency series reasonably well. The correlation of (the monthly aggregation of) our series and the monthly investment measure compiled by the Tax Authority is 0.70 when comparing year-on-year monthly growth rates. At the quarterly level, the correlation with the investment series from the national accounts (year-on-year growth rate) is much higher, 0.95.

We adjust our daily aggregate investment series for inflation similarly to how we handle sales and consumption: we apply the same monthly deflator to all daily observations in the month. To do this, we use the implicit price deflator obtained from monthly total investment series (nominal and real) provided by the Spanish Tax Authority.

2.1.4 Employment

For our measure of daily aggregate employment, we source a publicly available administrative series from the Spanish Ministry for Inclusion, Social Security and Migration, that records the total number of employment contracts registered in the Spanish Social Security system on any given day.

Enrollment in the Social Security system is mandatory for all employer-employee contracts in Spain, with the exception of a small number of private and public institutions which have historically maintained a separate pension system.¹⁵ In total, the series represent about 99% of the employed population. We should note here that a worker

¹⁵The largest of these, by assets, are the Spanish Bar Association and the Basque Autonomous Government.

may have more than one active contract in the system (e.g., someone maintaining two part-time jobs). Our series tracks the total number of active jobs registered in the Social Security system, rather than strictly the total number of employees.

The daily series is updated daily from Monday to Friday, and nets out job creation (new labor contracts registered with the social security system) from job destruction (labor contracts that have lapsed on that day) to obtain a daily series for the stock of employment contracts.¹⁶ This daily employment series starts from August 3rd, 2015.

Finally, at the monthly frequency, the Spanish Ministry for Inclusion, Social Security and Migration additionally provides a breakdown of monthly aggregate employment into aggregate permanent employment and aggregate fixed-term, or temporary employment, which we also use.

2.2 Other Data

As discussed above, we additionally compile more standard data on *(i)* interest rates, financial markets' and housing prices as well as consumer price indexes, both aggregate and by consumption category; *(ii)* a variety of expectations' data and confidence indicators; *(iii)* inventories; and *(iv)* a novel set of daily indicators of interest rate news coverage in Spanish media.¹⁷ We now briefly discuss each of these data.

Prices. For equities and housing we use, respectively, daily series data on stock prices from the IBEX35 index, and the monthly average transaction price of housing per square meter from the Statistical Information Center for Notaries (CIEN). The CIEN series is based on notarial records of purchase-sale contracts, and therefore reflects realized transaction prices rather than listing or asking prices.¹⁸ The daily data on stock prices begin on January 3rd, 2005, while the monthly data on transaction house prices begin in January 2007. Our measures of borrowing rates are the 1-, 3-, 6-, and 12-month Euribor rates—all available from January 4th, 1999—and the monthly home loan interest rates published by the Bank of Spain, available from June 2014.¹⁹ Finally, the series for the monthly consumer price index for Spain and COICOP-specific consumer price indexes both start in January 2012.

¹⁶For Saturdays and Sundays, we assume that the number of workers registered is the same as the previous Friday.

¹⁷Recall that we provide a summary of the data collected and respective sources in Appendix A.1.

¹⁸See Consejo General del Notariado, CIEN methodology documentation.

¹⁹We also collect data on credit quantities rather than prices. Monthly credit volumes for households and for firms are available from the Bank of Spain and start in June 2010.

Expectations and sentiment. We collect data on expectations for real activity, financial markets and prices from a variety of sources. We proxy for inflation expectations using market-based 1-, 5- and 10-year inflation-linked swaps for Spain, from Bloomberg. We gather consumer and business expectations on real activity in Spain from the European Commission’s harmonised Business and Consumer Surveys: industry, services, retail trade, construction, and consumer survey. Monthly Confidence Indicators (CIs), reflecting overall perception and expectations, are calculated separately for consumers and the four business sectors covered by the survey programme, and are available for Spain since January 2000. The monthly Economic Sentiment Indicator (ESI) and the Employment Expectations Indicator (EEI) are calculated based on a selection of questions from these surveys. We also make use of a recent consumer expectations survey, conducted by the European Central Bank since April 2020. This survey provides information on consumer expectations in Spain on various aspects of real activity and development in financial markets, such as expected mortgage interest rates and access to credit for the next 12 months.

Inventories. We measure inventory dynamics with monthly series from two sources. The first, the Spanish National Statistics Institute’s (INE) Short-term Stock and Inventory Survey, records the level of inventories held in wholesale and retail trade since January 2013. The second comprises the balances summarising how retail-trade and industry firms in Spain assess their stocks in the same European Commission harmonised Business and Consumer Surveys used above for our confidence indicators, seasonally adjusted and available since January 2000.²⁰

Interest rate news. In addition to financial market and survey-based expectation measures, we construct a novel set of daily indicators of interest rate news for Spain. These are based on the Global Database of Events, Language, and Tone (GDELT; Leetaru and Schrodtt, 2013), an open-source database that monitors digital, print, and broadcast news worldwide in more than 100 languages, updated every 15 minutes. From this database, we extract all news articles tagged with the theme INTEREST RATE and referencing Spain. We build three measures. First, a coverage indicator, defined as the share of articles mentioning interest rates over the total number of news articles recorded in Spain on that day. Second, a tone indicator, equal to the average sentiment score of those articles,

²⁰Survey responses are aggregated into balances, expressed in percentage points of total answers. For these three-option questions (stocks above normal, normal, or below normal), with P , E and M the respective shares of firms ($P + E + M = 100$), the balance is $B = P - M$. DG ECFIN compiles the national results and seasonally adjusts the resulting series.

as computed by GDELT’s dictionary-based natural language processing, multiplied by -1 so that negative values correspond to contractionary monetary policy news. Third, a combined Coverage–Sentiment (Overall News Intensity) index, constructed as the product of relative coverage and tone. These news-based indicators are available from 2017 to October 2023 and provide a direct measure of how quickly monetary policy shocks are disseminated through the news environment, offering a novel way to assess both the mechanism and the speed through which agents learn about monetary policy.

3 Methodology

We study the transmission of monetary policy combining the high-frequency data discussed in the previous section with a standard high-frequency approach to monetary policy shock identification. In this section, we review our baseline choices regarding shock identification, seasonal adjustment, and regression specification. Whenever possible, we opt for well-understood, off-the-shelf methods—we assess robustness in Section 7 of the paper.

3.1 Identification

We identify the effects of monetary policy shocks on the macroeconomy by using a standard series of monetary policy surprises for the Euro Area made publicly available by [Jarociński and Karadi \(2020\)](#). These daily monetary policy shocks are constructed from Euro Area high-frequency changes in financial assets around ECB policy announcements along the lines of [Gürkaynak, Sack and Swanson \(2005\)](#).

The shock series by [Jarociński and Karadi \(2020\)](#) addresses concerns with the relevance of a central bank “information channel”, potentially confounding monetary policy shocks. As argued by, for example, [Nakamura and Steinsson, 2018](#), [Cieslak and Schrimpf, 2019](#), and [Miranda-Agrippino and Ricco, 2021](#), information effects could mute the negative stock market response to changes in the expected policy path, or even reverse its sign. To control for the information channel, [Jarociński and Karadi \(2020\)](#) combine standard high-frequency identification with sign restrictions in a Bayesian VAR. Specifically, sign restrictions are applied to changes in the interest rate and stock market, where pure monetary policy shocks are associated with a negative comovement between interest rates and stock prices, while information shocks are associated with positive comove-

ment.²¹ The updated version of the monetary policy shocks database by Jarociński and Karadi (2020) that we use in our analysis includes 293 ECB policy announcements from 1999 to 2023, with 63 of them occurring during our baseline sample from August 2015 to October 2023. The standard deviation of the shock is 3.7 basis points in our consumption and employment samples and 4.1 basis points in the shorter sales and investment samples. Appendix B.1 presents further descriptive statistics for our baseline series of monetary policy shocks.²²

For robustness, in Appendix G.3 we consider two alternative series of monetary policy shocks using the Euro Area Monetary Policy Event-Study Database (Altavilla et al., 2019). The first series involves 1-month Overnight Indexed Swap (OIS) changes around policy announcements, controlling for the information channel by excluding instances in which the OIS rates and the Euro Stoxx 50 index stock price move in the same direction.²³ The second series is an updated version of the Policy Target factor by Altavilla et al. (2019), focusing on the short end of the yield curve and derived using factor analysis based on changes in the yield curve.

3.2 Seasonal Adjustment of Daily Time Series

Seasonally adjusting daily series faces at least three challenges. First, daily data are sensitive to calendar effects, such as the different number of working days. Second, one needs to purge not only predictable annual seasonal patterns, but also within-month and within-week regular variation. Third, noise is more pervasive in daily series, be it because true irregular variation in data is not time-averaged away or because measurement error may be heightened at high frequencies.

To deal with the first two challenges, we adopt a simple and transparent two-step approach to construct the dependent variable entering our local projections. In the first

²¹For the US, the literature has proposed a number of alternative ways to orthogonalize monetary policy surprises from the information channel. For example, Miranda-Agrippino and Ricco (2021) orthogonalize monetary surprises by projecting market-based monetary surprises on their own lags, and on the central bank’s information set, as summarized by Greenbook forecasts. More recently, Bauer and Swanson (2023) were able to produce monetary policy shocks that control for the information channel using publicly available data only.

²²Appendix B.3 tests a necessary condition for lead-lag exogeneity (Stock and Watson, 2018; Jordà and Taylor, 2025)—that the identified shock is not predictable from past outcomes—by extending the baseline local projection to negative horizons ($h = -60, \dots, -1$). We find no evidence of predictability: the fraction of pre-shock coefficients that reject zero at the 5% level is consistent with the nominal size of the test, and the impulse response emerges discontinuously at $h = 0$.

²³It is worth emphasizing that, while the identification strategy ensures that our shocks are orthogonal to the information channel, households naturally update their expectations once a shock occurs; this updating of expectations in response to the shock is a feature of transmission, not a sign that the response is contaminated by information effects.

step, we regress each raw daily series—in levels—on a full set of day-of-the-week and day-of-the-month fixed effects, cleaning out predictable within-week variation (e.g., systematically lower transactions on weekends) and within-month variation (e.g., payroll-driven spending spikes). In the second step, we compute year-on-year log growth rates on the adjusted series, differencing out any remaining annual periodicity. Throughout the paper, almost all daily series are seasonally adjusted using this baseline data transformation; the exception is financial-market variables, which remain in levels or log-levels and are not seasonally adjusted. As for the third challenge—daily noise—rather than pre-smoothing the data with a moving average, we address it at the estimation stage, by employing Smooth Local Projections described in the next subsection.

In Appendix G.2, we show that applying a range of alternative methods to deseasonalize and smooth daily data—including backward-looking moving averages coupled with standard LP, and model-based unobserved-components approaches—produces results that are qualitatively and quantitatively similar to those derived under our baseline.

3.3 Smooth Local Projections

We estimate daily impulse response functions (IRFs) to monetary policy shocks up to horizon H using the Smooth Local Projections (SLP) method of [Barnichon and Brownlees \(2019\)](#). The underlying specification is the standard local projection of [Jordà \(2005\)](#). For each horizon h , the impulse response coefficient $\beta_{h,0}$ is identified from:

$$y_{t+h} = \alpha_h + \sum_{\ell=0}^k \beta_{h,\ell} shock_{t-\ell} + \sum_{\ell=1}^p \varphi_{h,\ell} y_{t-\ell} + \theta_h cases_t + \delta_h stringency_t + \varepsilon_{h,t}, \quad (1)$$

where y_{t+h} is the dependent variable of interest, and $shock_t$ is the monetary policy shock at time t . At the daily or weekly frequency, we find no significant autocorrelation in the monetary policy shocks, and because of that, we do not include lags, $k = 0$. However, when monetary policy shocks are aggregated to monthly or quarterly frequency, we find significant shock autocorrelation that only dies out after 6 months. Hence, for monthly and quarterly series we include six and two lags of monetary policy shocks in our set of control variables, respectively.

Rather than estimating (1) horizon by horizon via OLS, SLP parameterizes the impulse response coefficients $\{\beta_{h,0}\}$ as a linear combination of B-spline basis functions and estimates the basis coefficients via penalized regression, imposing smoothness across horizons. The roughness penalty applies only to the shock coefficients; the coefficients on

control variables—lags of the dependent variable, COVID controls—are estimated freely at each horizon. This is particularly well-suited to our setting: daily impulse responses span hundreds of horizons, and unrestricted horizon-by-horizon OLS estimates can be excessively noisy. By smoothing the estimated IRF rather than the underlying data, SLP reduces the variance of the impulse response estimates while preserving the flexibility of the LP framework—and, importantly, avoids the need to pre-smooth the data with a moving average.²⁴ Following the penalized B-spline (P-spline) approach of [Barnichon and Brownlees \(2019\)](#), the roughness penalty is a second-order difference penalty ($r = 2$) on adjacent B-spline coefficients—shrinking the estimated impulse response toward a line—with its strength governed by a tuning parameter λ , which we select via five-fold cross-validation.²⁵

In the baseline specification, we estimate IRFs up to $H = 364$ days (one year after the shock) and include $p = 90$ lags of the endogenous variable (a quarter of past information).²⁶ Because our variables are measured in year-on-year growth rates, the 364-day horizon ensures that impulse responses up to one year can be directly interpreted as changes in levels; see [Appendix B.2](#) for further detail. The inclusion of lags of the dependent variable as controls is motivated by [Montiel Olea and Plagborg-Møller \(2021\)](#), who show that lag-augmenting local projections both render inference more robust and simplify standard error calculations, by avoiding the need to adjust for residual serial correlation.²⁷ We compute 68% and 90% confidence intervals using heteroskedasticity and autocorrelation consistent (HAC) standard errors.²⁸

²⁴[Barnichon and Brownlees \(2019\)](#) demonstrate, through simulations and empirical applications, that SLP provides more accurate and smoother IRF estimates than standard LP, especially when the data are noisy.

²⁵Specifically, the smoothing parameter λ is chosen from a grid of 18 candidate values ranging from 0.001 to 2. Cross-validation uses five contiguous time blocks—rather than random splits—to preserve the temporal dependence structure of the data. For each candidate λ and each fold, the model is estimated on the remaining four folds and the mean squared prediction error is computed on the held-out block. To avoid selecting an excessively small λ that overfits noise, we combine two selection rules: a marginal improvement criterion that identifies the first λ at which the relative reduction in cross-validation loss falls below a tolerance, and a curvature-based elbow rule. When the two rules agree, we select the more conservative (larger λ); when they disagree, we select the λ with lower cross-validation loss. For daily data, cross-validation is performed over a shorter horizon ($H = 90$) for computational tractability, and the selected λ is then used to estimate the full-horizon impulse response.

²⁶We report projections over a two-year horizon in [Appendix C.3](#), with the caveat that our comparatively short sample limits inference at longer horizons.

²⁷Combined with the year-on-year transformation, this choice can sharpen precision at horizons approaching one year: the base-year level that a long-horizon observation differences out, reappears as the recent endpoint of one of the included lags: this can raise the R^2 and narrow the confidence intervals there; [Appendix B.2](#) gives the details.

²⁸Results are unchanged when computing standard errors using alternative lag truncation choices or heteroskedasticity-robust standard errors without the autocorrelation correction.

Since our baseline sample includes the period of the COVID-19 pandemic, we adopt the approach advocated by [Schorfheide and Song \(2024\)](#) (also discussed in [Lenza and Primiceri, 2022](#)) and drop observations between March 14th, 2020 and October 30th, 2020 when estimating our baseline local projections.²⁹ In addition, we include two COVID controls: $cases_t$ is the log of new confirmed cases of COVID-19, and $stringency_t$ is the log of the stringency index.³⁰ Appendix [G.4](#) presents robustness checks using pre-COVID-19 sample periods, demonstrating that the full sample’s baseline method for addressing COVID-19 produces results comparable to those from samples ending before the eruption of the COVID-19 crisis.

For robustness, we consider two variations on our baseline procedure. First, in our baseline we use the identified monetary policy shocks of [Jarociński and Karadi \(2020\)](#) (which isolate the policy component of high-frequency surprises via sign restrictions) directly in the LP specification. The two alternative shock measures in Appendix [G.3](#) are the OIS 1-month and the Policy Target Factor—these are raw market-based surprises that do not strip out the information channel and are thus better viewed as instruments for monetary policy rather than identified shocks ([Stock and Watson, 2018](#)). We present estimates employing these two series in Appendix [G.6](#), showing that results are similar across LP and LP-IV specifications. Second, Appendix [G.2](#) considers alternative approaches to dealing with seasonality and noise in daily data, including applying a backward-looking moving average to the data and estimating standard (unsmoothed) LP via OLS, as well as model-based unobserved-components seasonal adjustment methods. Our baseline findings are robust to all these alternatives.

4 Slow vs. Fast Moving Variables in the Transmission of Monetary Policy: Novel Evidence at High Frequency

According to a widely held view in the community of central bankers, monetary impulses transmit to the economy by affecting financial markets quickly, on impact—namely, borrowing conditions, and stock prices—but only slowly affect production, in-

²⁹Alternatively, [Lenza and Primiceri \(2022\)](#) allow for time varying volatility in the residuals and obtain very high volatility estimates during COVID, so that effectively COVID observations contribute little to the estimation. Our strategy is more conservative, excluding any effect of extreme observations on our LP estimates altogether. Second, note that we do allow for COVID data to appear as lagged controls (for post-COVID observations for the dependent variable which we retain) in order to preserve the autocorrelation structure in data. We thank Giorgio Primiceri for discussions on this topic.

³⁰Covid cases data for Spain is compiled by the World Health Organization, and the stringency index is calculated by the University of Oxford Coronavirus Government Response Tracker.

vestment and consumption decisions by firms and households, with the corresponding macro aggregates responding only after long lags; see, for example, [Burr and Willems, 2024](#) for an institutional account of monetary policy transmission as a two-stage sequential process of fast vs. slow adjustment.³¹ In turn, this view resonates with the hypothesis in the academic literature, that some real variables can be expected to react more slowly than others. In the taxonomy of the seminal paper by [Bernanke, Boivin and Elias \(2005\)](#), the list of slow-moving variables spans employment, consumption, personal income, hourly earnings, CPI as well as industrial production. A large literature has built on this taxonomy, either contributing evidence, or using it for identification purposes. In this section, we leverage our novel high-frequency data to revisit it, producing granular evidence on the lags of the transmission of monetary policy across a broad set of financial and real variables.

4.1 Setting the Stage: Interest Rates and Inflation Expectations

To set the stage of our analysis, in Figure 1 we show the responses to a contractionary monetary shock of key interest rates, asset returns and asset-price-based inflation expectations—the variables defining the first stage of the transmission mechanism according to the consensus view. These IRFs—and all that follow—are calculated relative to one-standard-deviation contractionary monetary policy shock.


The daily responses of key interest rates, the 6- and 12-month Euribor, are shown in the first two panels (top row) of this figure.³² Both rates rise by roughly 1bp within the first 60 days following a monetary policy shock; the 12 month Euribor displays a continued rise thereafter. Correspondingly, the benchmark stock market index for Madrid’s stock exchange contracts—shown in the third panel (top row)—significantly on impact, by about 0.4%, remaining around that level for the rest of the year.

The second row of the figure shows the daily responses of inflation expectations using 1-, 5- and 10-year inflation-linked swaps for Spain. Consistent with a contractionary shock, these market-based inflation expectations decline in the first 30 days from the shock. While this decline happens across the term structure, it is quantitatively stronger for 1-year ahead expectations, with a drop of about 2bp.

³¹According to the Bank of England’s Quarterly Bulletin ([Burr and Willems, 2024](#)), the “Monetary Transmission Mechanism can be broken down into a first stage, the pass-through from Bank Rate to various asset prices and other interest rates, and a second stage – governing how financial conditions affect macroeconomic outcomes. First-stage transmission is typically rapid, thanks to the fast speed at which financial markets react to news, whereas the second stage takes more time.”

³²It is worth noting that these are the reference rates for most variable-rate mortgage contracts in Spain, themselves the most common contract in Spain.

Figure 1: Daily response of financial markets to a monetary policy shock



Notes: Financial market series are not seasonally adjusted. LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample starts in August 1st, 2015 and ends in October 30th, 2023. The monetary policy shock standard deviation is 3.7bp. Appendix D.2 reproduces this figure with the corresponding standard (unsmoothed) local projection overlaid.

To sum up, key interest rates, asset prices and expectations all move *fast*: they respond on impact, with economically and statistically significant reactions observed at high frequency, within the first few days and weeks after the monetary policy shock. Importantly, these results are not specific to Spain, our sample period or particular methodological choices. They are qualitatively and quantitatively consistent with a large literature documenting similarly fast and sizeable reactions of financial markets across different countries, time-periods and monetary policy shock measures;³³ Appendix C.1 further benchmarks our findings here, showing they are consistent with a broad range of published estimates.

³³See, for example: Jarociński and Karadi (2020) who documents monthly IRFs of the German one-year government bond yield and the Euro-Area stock index using a sample from January 1998 to December 2016, with a comparable shock magnitude of 3.5bp; Jarociński (2024) for their estimates of daily IRFs of financial variables in response to monetary policy shocks for the United States; Swanson (2021) for their IRF of 6-months US treasury yields; Miranda-Agrippino and Ricco (2021) for their IRFs of the 1 year T-bond and the S&P500; and Altavilla et al. (2019) for the impact of monetary policy shocks on 6-months and 1-year German yields.

4.2 Four Key Daily Measures of Economic Activity

Our baseline evidence on the effects of monetary policy on the real economy is presented in Figure 2. This figure shows the responses to a monetary policy shock of four daily measures of real economic activity: corporate sales (first row); two major components of aggregate demand, consumption (second row) and investment (third row); and employment (fourth row). The left column traces each response over a 365-day horizon and overlays two estimators—a standard, unsmoothed local projection on the seasonally adjusted data (in grey) and our smooth local-projection baseline (in red, with 68% and 90% confidence bands); the right column zooms in on the smooth baseline over the first 93 days.³⁴ To keep both estimators legible, the left column lets the vertical axis vary by row, while the zoom uses a common scale. The smooth response traces the center of the noisier raw estimates, so the smoothing step aids legibility rather than driving the result: for sales and consumption the short-lag contraction is already apparent in the raw series, and employment, measured from administrative population data, is estimated precisely throughout. Investment is the exception—its raw series is too volatile to read the response directly, so only the smooth estimator recovers a clear signal.


Turning to the behavior of each individual series, as shown in the top panel of Figure 2, we see that real corporate sales, our broadest measure of economic activity - encompassing final consumption, investment and intermediate input sales - respond rather quickly. Its negative raw response is detectable – significantly, at the one standard-deviation level – as early as four days after a monetary policy shock and more than half of the days (57% of the days) in the ensuing first month display significant negative adjustment.³⁵ In turn, the corresponding smoothed estimates turn persistently significant (at the 68% level) 20 days after the shock, reaching a local trough of -0.50% at day 51. The response stabilizes thereafter, before declining again roughly 200 days after the monetary policy shock.³⁶ Its global trough, -0.52%, is attained around day 330.

³⁴Because the smooth local projection is well motivated on its own terms—we estimate it on seasonally adjusted data with a cross-validated penalty, as detailed in Section 3.3—we display this raw overlay only here, for the four headline activity variables, where it is most useful to show directly that the short-lag signal does not hinge on smoothing. For the remaining daily figures—the financial-market responses of Figure 1 as well as the consumption-category, sectoral-sales, and upstream vs. downstream results that follow—the main text reports the smooth-LP baseline alone, to keep the figures legible, and Appendix D.2 reproduces each of them with the corresponding standard, unsmoothed local projection overlaid.

³⁵We classify each daily coefficient as significantly negative (upper bound of the pointwise 68% confidence band below zero) or significantly positive (lower bound above zero)—a one-standard-error, one-sided criterion under which a symmetric, zero-mean null would place roughly 16% of horizons in each category.

³⁶A similar pattern of initial contraction, partial recovery, and renewed decline appears in other settings: Slacalek, Tristani and Violante (2020) (Figure 3) for consumption in Germany and Spain at quarterly

Figure 2: Daily response of real activity to monetary policy shock



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The left column overlays, over a 365-day horizon, a standard (unsmoothed) local projection on the seasonally adjusted data (grey) and the smooth-LP baseline (red), with the vertical axis allowed to vary by row; the right column zooms in on the smooth-LP baseline over the first 93 days on a common vertical scale. The sample ends in October 30th, 2023. The sample starts one year after the sample start of each series reported in Table A1. For sales and investment—the shorter series—the monetary policy shock standard deviation is 4.1bp, while for consumption and employment it is 3.7bp.

While quantitatively smaller relative to sales, the response of aggregate consumption, shown in the second row of Figure 2, follows a broadly similar pattern. Its raw response is significant (and negative) already in the first week after a monetary policy shock and remains so for 80% of the days in the first month. Tracing out longer horizon response via the implied smooth LP, we see that household consumption contracts at short lags, stabilizes and then again falls at longer lags. Relative to corporate sales, the short-run trough in the smoothed response of consumption occurs at a slightly shorter horizon, 44 days after the monetary policy shock, and is about 20% smaller, at -0.39%. The response turns economically and statistically insignificant about 100 days after the monetary policy shock. The second spell of contraction is relatively shallow—attaining a (significant) local trough of -0.28% about 300 days after the shock.

The same broad pattern characterizes the response of our proxy for aggregate investment, displayed in the third row of Figure 2. Our smooth LP point estimates also display a statistically significant short-run contraction of -0.36% at the 34 day mark. Nonetheless, relative to the other series, the response of investment is noisier. The higher noise in this series may reflect the nature of investment spending—spikier than consumption and sales at the daily level—as well as higher measurement error in its construction. The contraction in investment is clearer and more persistent at longer horizons, starting at 8 months after the shock. The global trough is as deep as -1.46%, at the one-year horizon.

Finally, compared with sales, consumption and investment, the quantitative response of employment, shown in the bottom panel of the figure, is qualitatively and quantitatively different. It is smoother and much slower. It remains economically negligible at short horizons: three months after the monetary policy shock, the response of employment, while statistically significant, is only -0.11%. The response flattens out thereafter, resuming its decay at roughly 200 days after the shock. Its strongest response within the first year is precisely on day 365, with a cumulative decline of -0.22%.

To summarize the dynamic responses in a single interpretable metric, we additionally compute cumulative monetary multipliers, defined as the ratio of the cumulative response of each real-activity variable to the cumulative response of the 12-month Euribor; we report the results here and display the implied figures in Appendix C.2. Cumulative multipliers offer a compact way to gauge the overall strength and persistence of monetary transmission, taking into account that the policy rate itself responds persistently before returning toward baseline—an exact analogue to the cumulative fiscal multipliers of Mountford and Uhlig (2009) and Ramey (2016), recently applied to other contexts by

frequency, and Miranda-Agrippino and Ricco (2021) (Figure 12) for U.S. industrial production in monthly data.

Barnichon and Mesters (2024), Fieldhouse, Mertens and Ravn (2018) and Ishan, Ramey and Klenow (2024).

At short horizons, the cumulative multiplier is relatively large in magnitude, reflecting the fact that demand contracts immediately after the monetary policy shock, while interest rates, still rising, do so only gradually. As policy rates continue to increase and reach their peak, the pace of the decline in demand and real activity stabilizes, producing a temporary flattening of the cumulative ratio. At longer horizons, interest rates begin to return toward baseline, while our macro variables continue to weaken. The cumulative multiplier once again becomes increasingly negative—especially so for employment.

Quantitatively, investment displays the largest cumulative multiplier, reaching approximately -0.5 after one year. Sales lie in between, reaching approximately -0.3 , while consumption reaches around -0.2 . Employment is negative throughout but comparatively small, with a cumulative multiplier of approximately -0.1 after one year, consistent with slower adjustment of labor inputs relative to demand. Figure C1 in Appendix C.2 reports the cumulative multipliers at each horizon up to one year for all four variables.

Taken together, these dynamics illustrate how the timing of monetary tightening relative to real adjustment shapes the overall strength of transmission. Because policy rate increases are themselves persistent, the area under the interest rate response expands even as real activity falls, making the ratio of cumulative responses a natural and transparent measure of the dynamic potency of monetary policy.

4.3 Taking Stock: Slow-moving Variables vs. the Slow Transmission of Short Lags

The key fact unveiled in this section is that monetary policy has economically and statistically significant effects on real economic activity already within weeks from a policy innovation. Major components of demand and gross output, typically considered slow in reacting, in fact closely track the “fast” response of asset prices and expectations to a monetary policy shock. Our evidence thus casts a different light on the consensus view of monetary transmission, as well as on popular labelling of variables into slow- and fast-moving used in identification.

At the same time, Friedman’s dictum remains visible in employment and prices. As we have seen, the response of employment is smooth and becomes economically significant at longer lags—when sales, consumption and investment also align in a contraction. The employment response reaches its trough only around one year or later. In Section 7.3,

using monthly data, we will show that the response of employment is actually mirrored by CPI inflation; see Figure 13 therein. Arguably, employment and the CPI are two key variables for central banks: our evidence shows that monetary policy drives them down in tandem, along with the other three aggregates in Figure 2, at longer lags.³⁷

What our evidence qualifies is the interpretation of these long lags as reflecting a generic “slow response of real variables”. The contraction in demand is instead already economically and statistically significant at short lags, reaching its global trough early. The longer-lag comovement of economic activity, employment and, eventually, inflation therefore points instead to mechanisms that slow down the propagation of this initial contraction through the economy: from demand at short lags, to employment and broader economic activity at longer lags, and eventually to prices.

5 Monetary Transmission Across Goods and Sectors

In this section, we exploit the rich information in our dataset to study the heterogeneity in the short-lag response across subcategories of demand and sales. In particular, we are interested in establishing whether, underlying our aggregate responses, there are significant and economically meaningful differences in the response to a monetary shock across (i) subcategories of final consumption demand for goods, contrasting differences in the response of durables and luxuries vs. that of non-durables and necessities; (ii) subcategories of corporate sales, where we further distinguish between downstream sectors (those closer to final household demand) and upstream sectors (those further away from final household demand).

5.1 Response Lags Across Goods Categories: Final Demand for Durables and Luxuries Responds Fastest

Our data on consumption is dense enough in the cross-section to allow finer, daily disaggregated cuts that are typically unavailable at high-frequency. Specifically, we are able to decompose our aggregate consumption series at the two-digit COICOP category

³⁷In Appendix C.3 we extend the analysis to a two-year horizon, where the employment response is already near its trough around one year, while the CPI reaches its trough only around 20 months. Whether these long-lag troughs lie at around one year, at one to two years (Christiano, Eichenbaum and Evans, 1999), or at two to three years (Romer and Romer, 2004), is an open question our short sample cannot settle. Pinning down their precise location requires longer samples and methods built for long-horizon estimation. Our contribution is instead the much earlier global trough in final demand. Earlier studies’ later troughs may themselves partly reflect lower data frequency and the time-aggregation biases we document in Section 6.

level, obtaining eleven consumption series; see Appendix A.2 for a listing of the disaggregated consumption categories.

The results of our disaggregated analysis are shown in Figure 3. The categories of consumption goods that fall significantly in the immediate aftermath of a contractionary monetary policy shock are durable or semi-durable goods and luxury goods. The durable or semi-durable goods include clothing and footwear, transport³⁸ and, to a lesser extent, furnishing, equipment and maintenance. The luxury goods include restaurants and hotels, recreation and culture, health and education.³⁹ In contrast, the response of necessities—including food and non-alcoholic beverages, housing services and utilities and communication—remains subdued at all horizons within the first year after a monetary policy shock.⁴⁰ The upshot of these heterogeneous responses is that, on average, durable and luxury goods and services decline between 1% and 2% within the first 60 days from the shock. The response of essential goods instead ranges between 0% and 0.25%. In fact, the consumption of food and non-alcoholic beverages displays a slight rise, consistent with substitution away from restaurants and hotels (i.e., from “food away from home”).

Overall, these findings provide a detailed high-frequency counterpart to the literature documenting that the response of consumption to monetary policy shocks is disproportionately driven by the response of durable goods consumption—see, e.g., [Erceg and Levin \(2006\)](#), [Monacelli \(2009\)](#), [Sterk and Tenreyro \(2018\)](#), and [McKay and Wieland \(2021\)](#).⁴¹ We show that the contraction in the demand for these goods is not only sharp,


³⁸We classify the transport category as semi-durable because it includes both durable vehicle purchases and a mix of expenditures on transportation services—e.g., taxi rides or monthly public transportation pass—that are either non-durable (the service flow of a taxi ride extinguished within the day) or semi-durable (the service flow of a transportation pass lasts typically for a month). Using an independent monthly series specific to sales of vehicles, in Appendix F.3 we show that its response is very similar to the response of consumption of “Transport” services.

³⁹In Spain, the large majority of education and health services are publicly provided by the state at low or no cost to the end user. The transactions related to education in our data are mostly generated by spending on private education and private health services, which in Spain can be considered luxury goods.

⁴⁰We should note that the criteria used in constructing housing services and utilities and communication consumption may weigh on these findings. This is because, for these series, the daily consumption is computed by either imputing rents or distributing the monthly/bi-monthly utility payments over the days of the month/months. Two comments are in order. First, while the data construction procedure could create an artificial delay in the response, it would not prevent our model from eventually detecting a significant response. Second, the finding that, for housing services and utilities and communication, the responses are very small, can be cross-checked with the response obtained from daily sales data (for the information and communication sector), that are not constructed following the same criteria. The fact that the empirical response of sales is also very small, suggests that our results are not determined by the data construction criteria.

⁴¹The point is well-understood in the central bank community, as exemplified by the following quote: “Our staff estimate(s) that spending on durable goods and luxury goods declines more than it does on

Figure 3: Daily response of consumption by category to a monetary policy shock



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample starts in August 1st, 2016 and ends in October 30th, 2023. The monetary policy shock standard deviation is 3.7bp. See [Buda et al. \(2022\)](#) and Appendix A for further details on how the consumption categories were constructed and on their cross section and time series characterization. Appendix D.2 reproduces this figure with the standard (unsmoothed) local projection on the seasonally adjusted data overlaid.


but also fast—it is significant at very short lags.

5.2 Response Lags Across Sectors: Downstream, Final-demand Sectors Respond Fastest

We can further inspect the heterogeneous transmission of monetary policy shocks using the 20 disaggregated daily sectoral sales series made available by the Spanish Tax Authority; see Appendix A.2 for a full listing and associated descriptive statistics.

non-durables, services, and necessity goods” ([Lane, 2024](#)).

Figure 4: Daily response of sales by sector to a monetary policy shock



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample starts in July 1st, 2018 and ends in October 30th, 2023. The monetary policy shock standard deviation is 4.1bp. See Appendix A for further details on sales sectoral classification. Appendix D.2 reproduces this figure with the standard (unsmoothed) local projection on the seasonally adjusted data overlaid.

We start by considering a subset of nine sectors that best match the consumption categories studied above. The results, shown in Figure 4, suggest that the responses to a monetary shock for sectors that are closer to the final demand by households are broadly consistent with the short-lag response of consumption⁴²—when not slightly stronger. Specifically, wholesale and retail broadly track the response of aggregate consumption in Figure 2, with a decline of about 1% in the first 60 days of a monetary policy shock. The response of textile manufacturing mirrors the adjustment in consumption of clothing

⁴²It is worth stressing that this is so, despite differences in data sources and coverage between our consumption measures (being computed from banking transactions by private households) and firms' sales (measures from VAT tax declarations).

and footwear in Figure 3, with a marked, deeper trough during the second post-shock month. Other sectors where durable goods are highly represented in sales, such as electronics manufacturing or construction, also display patterns similar to the disaggregated consumption series, if only a bit noisier, with negative responses only in the short run. Finally, sales of the food manufacturing sector or information and communication display economically small responses, consistent with the patterns observed for the consumption of food and non-alcoholic beverages or communications expenses, respectively. In contrast, the responses of energy, transportation and storage or professional and administrative services are slower, but deeper and more protracted, with troughs at roughly 300 days from the monetary policy shock.

Importantly, our disaggregated results suggest that the responses of upstream sectors, providing general purpose inputs for the production of goods and services, may be different from downstream sectors. We submit this observation to a formal test. To do so, we return to the full daily disaggregated sales data for 20 sectors and bridge the Spanish Tax Authority sales sectoral classification with the INE Input-Output matrix to compute an upstreamness indicator following Antràs et al. (2012); Appendix D.1 provides details on how we bridge the sectoral classification of the Spanish Tax Authority with the INE IO sectoral classification, and on how we derive our upstream vs. downstream sectoral classification of the sales data. Based on this indicator, we classify a sector as upstream—far from final demand by households—or downstream depending on whether it scores above or below the empirical average of the indicator for all sectors.⁴³ As shown in the Appendix, in our classification, overall energy, transportation and storage, and professional and administrative services are accurately considered upstream sectors.


We then estimate the following panel Local Projection:

$$y_{t+h,s} = \alpha_{h,s} + \sum_{\ell=0}^k \beta_{h,\ell} shock_{t-\ell} + \sum_{\ell=0}^k \gamma_{h,\ell} shock_{t-\ell} \times up_s + \sum_{\ell=1}^p \varphi_{h,\ell} y_{t-\ell,s} + \theta_h cases_t + \delta_h stringency_t + \varepsilon_{h,t}, \quad (2)$$

where $\alpha_{h,s}$ is the sector fixed-effect and up_s is a dummy variable that takes value one if the sector is classified as upstream and zero if it is classified as downstream. This equation is the panel version of equation (1), with an added interaction term of the monetary policy shock using the upstream dummy. We estimate it with fixed-effects: $\hat{\beta}_{h,0}$ gives us the

⁴³Our results were consistent under an alternative sector classification, considering upstream if above the third quartile and downstream if below the first quartile.

Figure 5: Upstream vs. downstream sectoral sales to a monetary policy shock



Notes: Left panel displays the LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity-robust standard errors, clustered by time. Lighter-shaded areas are the 90% confidence intervals. The sample starts in July 1st, 2018 and ends in October 30th, 2023. The monetary policy shock standard deviation is 4.1bp. See Appendix D.1 for further details on sales upstream (red line and band) vs. downstream (dashed blue line and blue band) sectoral classification. Right panel presents the difference between upstream and downstream sales responses together with the 68% (lighter-shaded areas) and 90% (darker-shaded areas) confidence intervals of this difference. Appendix D.2 reproduces this figure with the standard (unsmoothed) local projection on the seasonally adjusted panel overlaid.

estimated response of sales at horizon h for downstream sectors, while $\hat{\beta}_{h,0} + \hat{\gamma}_{h,0}$ gives us the estimated response of sales at horizon h for upstream sectors.

Results are shown in Figure 5. The left-hand panel shows that the response lag of downstream sales—closer to final consumption demand—is short, statistically significant in about a month, and, compared to upstream sales, much stronger on impact.⁴⁴ The short-run trough of downstream sectors is almost three times deeper relative to the upstream sectors. Mirroring the pattern in the response of consumption that we have documented previously, downstream sales stabilize in the second quarter, and then contract further at long lags. Conversely, the response lag of upstream sectors is somewhat slower—it becomes statistically significant 60 days after the shock. The upstream response is, however, very persistent and the size of the contraction remains stable at longer lags. Finally, the right-hand panel of the figure plots the difference between the two groups of sectors, highlighting the statistically significant deeper short-run response

⁴⁴We explored the monthly sales data beginning in 2000 and confirmed that, in the pre-COVID sample ending in 2019, downstream sectors exhibit stronger contemporaneous responses to monetary policy shocks than upstream sectors. However, differences in sectoral classification between the monthly and daily sales data constrain the comparability of this analysis.

of the downstream part of the economy during the first quarter after the monetary policy shock, with the two series aligning over longer—6 months and above—lags.

Overall, these results provide a distinctive perspective on the short-term transmission of monetary policy. The contraction in upstream industries lags behind the strong and rapid response observed in downstream industries, suggesting that an initial adjustment in downstream final demand is transmitted gradually upstream following a monetary policy shock. At the same time, once the contraction begins, the decline in upstream industries’ sales appears to be more persistent. Taken together, these significant differences in the responses across industries highlight the potential value of explicitly incorporating production networks as a key determinant of the transmission of monetary policy shocks; see, for example, [Ozdagli and Weber \(2017\)](#) and [Ghassibe \(2021\)](#) for early empirical findings along these lines.

6 Time Aggregation

Unlike the high-frequency data we assemble in this paper, empirical measurements of continuous aggregate processes are typically available only at relatively low frequencies, at monthly or quarterly intervals at best. This, in turn, raises the well-known possibility of ‘time aggregation biases,’ whereby infrequent measurement may compromise the validity of empirical findings.⁴⁵ For example, in one of the first empirical assessments of the importance of such biases, [Christiano and Eichenbaum \(1987\)](#) already conclude that “temporal aggregation bias can be quantitatively important in the sense of significantly distorting inference” regarding the dynamic relation between money and output.⁴⁶ These issues have regained prominence in the modern literature relying on high-frequency identification of shocks, as these shocks must typically be “time aggregated” to match the lower frequency at which economic series of interest are available. As [Gertler and Karadi \(2015\)](#) note, simply summing high-frequency surprises over a calendar month can be problematic if shocks occur near the end of the period, as their full macroeconomic impact is delayed. However, [Ramey \(2016\)](#) warns that Gertler and Karadi’s proposed solution — constructing a monthly average of a rolling 30-day cumulative daily sum — mechanically introduces predictability and serial correlation into the instrument series.

In this section, we leverage our high-frequency dataset and our ability to line up the frequency of shocks with that of data to shed empirical light on potential issues

⁴⁵This is a classical question in economics and econometrics; see, e.g., the seminal contributions of [Amemiya and Wu \(1972\)](#), [Sims \(1971\)](#) and [Geweke \(1978\)](#), [Marcet \(1991\)](#) among others.

⁴⁶See also the corresponding comment by [Stock \(1987\)](#).

arising from time aggregation of data. We propose to do so by comparing two sets of estimates. The first is obtained by time-aggregating —by averaging— our daily local projection estimates at the weekly, monthly and quarterly horizons.⁴⁷ The second is obtained from local projections based on time-aggregated versions of our daily data—at weekly, monthly and quarterly frequency—where we sum the high frequency monetary policy disturbances within the period. That is, we compare time-averaged daily *estimates* with standard estimates obtained with time-aggregated *data*. As noted above (and in Section 3.3), autocorrelation may arise in monetary policy surprises following their aggregation to lower frequencies. Consequently, we incorporate six lags for monthly frequencies and two for quarterly frequencies of monetary policy shocks into our control variables. For these lower-frequency estimates, we apply standard local projections, estimated horizon by horizon via OLS, with the same treatment of the pandemic as in our baseline: we exclude the COVID window and include the COVID controls of Section 3.3. The aggregated series are the daily series seasonally adjusted in the simple manner described in Section 3.2, aggregated to each frequency.


6.1 Short lags at weekly and monthly frequency, but not at quarterly frequency

For our two new series, averaged daily impulse responses and direct lower-frequency impulse responses estimates, Figures 6, 7 and 8 plot the implied responses of sales, consumption, investment and employment, at the weekly, monthly and quarterly frequencies, respectively.

As shown in Figures 6 and 7, for all variables, the point estimates of the response using our aggregated daily series and using weekly or monthly aggregated data are close to each other—the confidence intervals nearly overlap. Patterns are similar in the two estimates: the responses of consumption and sales are significant in the first two months of a monetary policy shock; the short-run contraction is noisier for investment, but aligns over a large portion of the year horizon; the responses of employment is economically smaller, but quite persistent and significant over the horizon of our analysis. In our data,

⁴⁷Instead of displaying all IRFs at low(er)-frequencies, in Appendix E.1 we present an alternative approach where we plot the same low-frequency LP IRFs at a daily resolution; that is we plot, say, a LP based on aggregated monthly data against a daily axis, allowing us to superimpose our daily baseline local projection estimates (rather than averaging these to the monthly frequency as in main text). Comparisons may be harder visually. Yet, this alternative method highlights the finer details of high-frequency data, especially in investment, where responses become significant around 45 days after shock. It also highlights that sales, consumption and employment respond primarily at the end of the month, reinforcing the value of daily data granularity.

Figure 6: Time aggregation: Weekly responses




Notes: LP impulse response functions to a one standard deviation monetary policy shock. The shock size is identical across variables and frequencies: all responses are scaled by the standard deviation of the daily shock series, as in our baseline. The responses are reported in levels. The weekly estimates are standard (unsmoothed) local projections with the COVID treatment of Section 3.3, on weekly aggregates of the seasonally adjusted daily series of Section 3.2. Aggregated-data LP confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors; averaged daily LP intervals average the daily standard errors within each week. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in October 30th, 2023. The sample starts one year after the sample start of each series reported in Table A1. Clockwise, we display local projections for total sales, employment, consumption and investment. Dashed lines are the implied low-frequency aggregated data (weekly) LP point estimates, while solid lines are the weekly averages of daily LP estimates in our baseline.

time aggregation is not an issue when employing monthly or higher frequency series.

However, time aggregation does make a difference at the quarterly frequency—as shown in Figure 8. Using quarterly data, we find no significant same-quarter responses of sales and consumption, while we do so when using averaged daily responses at a quarterly frequency. In our data, quarterly aggregation shifts information to lower frequencies—responses become statistically significant only in the second quarter after the shock. It thus alters both the impact response to the shock and the overall time profile of the implied local projections during the first year after a shock.

Figure 7: Time aggregation: Monthly responses




Notes: LP impulse response functions to a one standard deviation monetary policy shock. The shock size is identical across variables and frequencies: all responses are scaled by the standard deviation of the daily shock series, as in our baseline. The responses are reported in levels. The monthly estimates are standard (unsmoothed) local projections with the COVID treatment of Section 3.3, on monthly aggregates of the seasonally adjusted daily series of Section 3.2. Aggregated-data LP confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors; averaged daily LP intervals average the daily standard errors within each month. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in October 2023. The sample starts one year after the sample start of each series reported in Table A1. Clockwise, we display local projections for total sales, employment, consumption and investment. Dashed lines are the implied low-frequency aggregated data (monthly) LP point estimates, while solid lines are the monthly averages of daily LP estimates in our baseline.

6.2 Time aggregation bias and its sources: simulation evidence

These results suggest that estimating local projections on daily data aggregated to sufficiently low-frequency time units produces a time aggregation bias. But without a concrete data-generating process, one cannot precisely define the bias, quantify its magnitude, or examine potential sources. To make progress on this problem, we specify a Monte Carlo simulation environment calibrated to key features of our data. The full description and results are in Appendix E, and here we provide a summary.

For concreteness, we focus on a DGP for daily consumption growth g_t and specify

Figure 8: Time aggregation: Quarterly responses



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The shock size is identical across variables and frequencies: all responses are scaled by the standard deviation of the daily shock series, as in our baseline. The responses are reported in levels. The quarterly estimates are standard (unsmoothed) local projections with the COVID treatment of Section 3.3, on quarterly aggregates of the seasonally adjusted daily series of Section 3.2. Aggregated-data LP confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors; averaged daily LP intervals average the daily standard errors within each quarter. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in October 30th, 2023. The sample starts one year after the sample start of each series reported in Table A1. Clockwise, we display local projections for total sales, employment, consumption and investment. Dashed lines are the implied low-frequency aggregated data (quarterly) LP point estimates, while solid lines are the quarterly averages of daily LP estimates in our baseline.

the AR(1) process to accommodate intrinsic persistence in consumption growth:

$$g_t = \alpha + \rho g_{t-1} + \sum_{h=0}^{365} \theta_h s_{t-h} + \gamma^T \mathbf{X}_t + \epsilon_t. \quad (3)$$

Here s_{t-h} is the monetary policy shock on day $t-h$, $\mathbf{X}_t$ are the COVID controls in our baseline local projections, and $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$ is a white-noise disturbance term. We set the θ_h coefficients so that the IRF produced by (3) is equivalent at each horizon to the estimated smooth IRF for consumption in Figure 2. We then fit the process to our actual consumption growth data, using the monetary policy shocks observed in our sample, to estimate the remaining parameters.

Once parameterized, we draw 1,000 simulations of daily consumption growth. Specifically, for each simulation, we (i) draw a sequence of monetary policy shocks calibrated to the empirical surprises;⁴⁸ (ii) draw consumption growth from (3); (iii) convert the simulated growth series to implied daily levels, aggregate to quarterly frequency, and compute annual growth rates; (iv) fit a standard local projection on quarterly data using the same specification as in (1). Time aggregation bias at horizon h is the difference between the average point estimate across simulations and the true IRF at that horizon.⁴⁹

Table 1: Quarterly IRF Estimation: Bias and Coverage Across Simulation Designs

Sample	IRF	Shocks / Q	Q0			Q1		
			Bias	68%	90%	Bias	68%	90%
IRF in DGP (Consumption): Q0 = −0.281, Q1 = −0.061								
Baseline	Cons	Two	+0.167	10.2	23.8	−0.124	19.7	36.0
Extended	Cons	Two	+0.153	0.4	1.7	−0.134	0.1	0.7
Extended	Cons	One (First-day)	−0.006	61.9	83.7	−0.001	65.3	87.6
Extended	Cons	One (Early)	+0.056	23.5	43.4	−0.051	26.2	47.3
Extended	Cons	One (Late)	+0.246	0.0	0.0	−0.221	0.0	0.0
IRF in DGP (Employment): Q0 = −0.082, Q1 = −0.105								
Baseline	Emp	Two	+0.059	38.3	61.8	+0.004	64.9	84.8
Baseline	Emp	One (First-day)	+0.001	61.4	84.1	+0.001	60.5	80.8

Notes: We design a variety of simulations to study time aggregation bias and draw 1,000 datasets in each. Bias is the mean point estimate minus the value of the IRF in the data generating process (DGP). Coverage is the percent of simulations where the given confidence interval contains the true value. The baseline (extended) sample runs from October 29, 2016 through June 15, 2023 (December 31, 2043). The Consumption and Employment IRFs are those estimated from real data and displayed in Figure 2. In the real data, there are two non-zero shocks per quarter which we replicate in the ‘Two’ designs. ‘One (First-day)’, ‘One (Early)’ and ‘One (Late)’ are alternative designs where data is simulated using a single shock each quarter, placed on the first calendar day of the quarter, at the position of the first observed shock, or at the position of the second observed shock, respectively. All reported estimates are from standard local projections estimated on quarterly-aggregated data.

The first row of Table 1 presents quarterly aggregation results for the first two quarters, i.e. impact (Q0) and first quarter post-shock (Q1).⁵⁰ The true value of the consumption IRF at Q0 is -0.281 while the average point estimate across simulations is -0.114,

⁴⁸We specify a student-t distribution for the shocks whose parameters are set to match the mean and variance of the policy surprises in our sample. To simulate a sequence of monetary policy shocks, we draw independently from this distribution for each date with an ECB meeting in our sample, and set $s_t = 0$ for all other dates.

⁴⁹The true value of the IRF is the average value of the smoothed LP coefficients for consumption from Figure 2 within the reference period. For example, the true value for Q0 is the average over days 0 through 90.

⁵⁰Appendix E reports results for all horizons. Bias and coverage distortions appear at longer horizons but generally improve by Q3.

which implies a bias of 0.167. We also report coverage of the 68% and 90% confidence intervals. Recall that coverage is the fraction of simulations in which the respective intervals contain the true value and is arguably the relevant quantity for inference. These coverage ratios for each confidence level are 10.2 and 23.8, respectively, well below nominal values. Q1 instead features a large negative bias of -0.124 and also poor coverage. These findings replicate the qualitative results for consumption from Figure 8 and show they indeed result from time aggregation bias.

What generates time-aggregation bias? While a full theory is outside the scope of the paper, variations of the baseline simulation can provide insights. To examine whether the bias arises from the short quarterly sample, we extend the simulation through December 31, 2043, increasing the number of quarters from 25 to 108, but otherwise keeping the data-generating process as in the baseline. This has little effect on bias but substantially worsens what was already poor coverage (Table 1, second row). This is because the longer series act to shorten confidence intervals at wrongly centered values.

Next, we consider the timing of shocks.⁵¹ For simplicity, suppose each quarter has 90 days and a single shock that arrives D days before the quarter ends. A local projection fit on quarterly data will, at its earliest horizon, capture the behavior of the daily IRF averaged over its first D days. The remaining $90 - D$ days of the daily IRF response over its first quarter will be estimated as part of the second-quarter response. To examine estimates *without* this effect, we alter the simulation so that a single shock arrives per quarter on the first day of the quarter. As Table 1 (third row) shows, this essentially eliminates bias and restores near-nominal coverage.

To explore more realistic timing, note that each quarter in our data has exactly two non-zero shocks. The average distance between the first (second) shock and the end of the quarter is 68 (20) days. Recall that the consumption IRF has a short-run trough at 44 days that dissipates after roughly 100 days. The late shock is therefore likely to drive much of the time aggregation bias. To test this, we use two alternative simulations, each with a single shock per quarter, placed near the first and the second shock in the actual data, respectively. Bias for the early shock is moderate, although not fully eliminated. For the late shock, the estimated impact effect is almost zero with zero coverage (Table 1, rows four/five).

Finally, we ask whether the shape of the consumption IRF is responsible for generating time aggregation bias. The IRF for employment provides a natural contrast, as it is relatively slow moving and monotonic. We alter the θ_h coefficients in (3) to reflect the

⁵¹Shock timing is emphasized in Gertler and Karadi (2015) (see footnote 11) and more recently in Jacobson, Matthes and Walker (2026) in the context of the price puzzle.

estimated smooth IRF for employment from Figure 2, and repeat the baseline simulation. There continues to be a bias at Q0 that is actually *worse* than the baseline in percentage terms, although coverage improves since absolute bias is smaller. While in Q1 the bias largely disappears, Table E5 shows it reappears in Q2 and Q3. Thus, time aggregation bias is not unique to a particular IRF shape.

In summary, we document that in situations where researchers can only fit local projections on low-frequency data that average over unseen daily dynamics, a time aggregation bias emerges whenever the timing of shocks does not align with the timing of the start of each reference period. This highlights the importance of having a full system of daily time series to pinpoint the unfolding of effects following a shock.

7 Extensions and Robustness

In this section, we combine our daily data with a larger set of variables available at the monthly, but not higher, frequency—relying on our time aggregation result, that short lags are preserved when going from daily to monthly. The analysis proceeds in three parts. We first ask how monetary policy reaches households’ final demand so quickly, documenting the transmission of policy decisions through news coverage, credit conditions and expectations, and the differential response of consumption across households more or less exposed to credit markets. We then turn to what slows the passage from final demand to production, showing that inventories buffer the gap between fast-moving sales and slow-moving output, and that employment and prices adjust at longer lags. We conclude by reporting a series of robustness exercises, varying in turn the data sources, the treatment of seasonality and noise, the measure of the monetary policy shock, and the sample period.

7.1 The Fast Transmission to Households: Information Flows, Credit Conditions, and Expectations

The financial market responses documented in Figure 1—the immediate rise in short-term rates, the contraction in equity valuations, and the decline in asset-price-based inflation expectations—set the stage for the household responses we document in our baseline. In this subsection we ask how, and how quickly, these adjustments reach households and feed through to consumption. We bring evidence on three complementary channels: the rapid diffusion of interest rate news through the media; the fast adjustment of credit conditions and asset prices that directly affect household balance sheets; and the simulta-

neous deterioration of broader macroeconomic expectations. We then use cross-sectional variation across households in their direct exposure to credit markets to exemplify how fast information updates translate into fast spending decisions.

To assess how quickly households can learn about a monetary policy shock, we construct a novel set of daily indicators of interest rate news for Spain using the Global Database of Events, Language, and Tone (GDELT; [Leetaru and Schrod, 2013](#)). GDELT monitors digital, print, and broadcast media worldwide in more than 100 languages, updated every 15 minutes; from this database we extract all articles from 419 Spanish media outlets—including national and regional newspapers, digital portals, and broadcast sources—that are tagged with the theme `INTEREST_RATE`, a classification based on a pre-defined set of keywords including “interest rate”, “policy rate”, “rate hike”, “rate cut”, “monetary policy”, “tightening”, and “easing”.⁵² We build three daily measures over the period January 2017–October 2023, which overlaps with our baseline sample: a coverage indicator (the share of all Spanish articles on GDELT that mention interest rates on a given day), a tone indicator (the average sentiment score of those articles, with more negative values indicating more negatively toned coverage of interest rates), and the combined Overall News Intensity index, which captures the simultaneous extent and negativity of interest rate news by multiplying coverage and tone.⁵³

Figure 9 shows the daily IRFs of all three indicators to a contractionary monetary policy shock. Coverage of interest rate news rises by about 0.02 percentage points on impact—roughly 3% of its daily mean—and remains significantly elevated for roughly 30 days before returning to baseline. Tone deteriorates on impact by about 0.02 index points.⁵⁴ The Overall News Intensity index falls significantly on impact—by about 8% of its absolute mean—reflecting the joint spike in coverage and increased negativity of tone. Together, the three series show that monetary policy surprises are disseminated almost immediately through the media: interest rate coverage jumps and its tone becomes more negative within days of the shock, giving households a fast signal that borrowing conditions have tightened.


The actual financial developments behind these news are documented in Figure 10,

⁵²The Baker-Bloom-Davis EPU index for Spain and the Ghirelli, Pérez and Urtasun (2019) index are constructed at monthly frequency. The Caldara-Iacoviello (2022) Geopolitical Risk index has a Spain series, but relies solely on coverage in the NYT, FT, and WSJ; our indicators are daily and based on the domestic media environment Spanish households actually face.

⁵³Appendix F.3 provides further details on data construction and descriptive statistics. The coverage indicator has a daily mean of 0.76% (standard deviation 0.63%), reflecting the share of interest-rate articles in normal times; the tone indicator has a mean of -1.31 (standard deviation 0.83) on a scale roughly bounded by $[-10, +10]$, so that interest rate news is, on average, negatively toned.

⁵⁴[Ashwin, Kalamara and Saiz \(2024\)](#) finds that generic sentiment dictionaries, which GDELT implements, outperform economically targeted dictionaries in nowcasting euro area growth.

Figure 9: Daily response of interest rate news coverage and sentiment to a monetary policy shock




Notes: LP impulse response functions of our three news-based indicators—coverage, tone, and the Overall News Intensity index—to a one standard deviation contractionary monetary policy shock. The coverage indicator is the share of Spanish articles on GDELT tagged with the theme INTEREST_RATE over all articles recorded in Spain on that day; the tone indicator is the average sentiment score of those articles as computed by GDELT’s dictionary-based natural language processing, with lower values indicating more negatively toned coverage; and the Overall News Intensity index is their product. GDELT monitors digital, print, and broadcast media worldwide; our Spanish extract covers a large number of domestic and international outlets reporting on Spain. Confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Sample: January 2017–October 2023. Monetary policy shock SD: 3.7bp.

which shows the monthly IRFs of credit volumes, mortgage rates, and house prices. Credit to households declines by about 0.37% already in the first month after the monetary policy shock, attaining a deeper trough of about -0.79% at 10 months. Credit to firms, by contrast, contracts significantly only from the second quarter on, with a decline of -0.4% thereafter; the impact response is a mild expansion, possibly reflecting temporarily higher short-term borrowing needs. House prices drop significantly from the third month on, reaching a trough of -0.35% at month 6.⁵⁵ Mortgage rates rise on impact, with a peak response of 1bp, and remain elevated for six months before turning insignificant.

Faced with news coverage and the financial effects of monetary shocks, households and firms update their expectations about broader macroeconomic prospects—including future income, employment outlook, and general economic conditions. Figure 11 reports the monthly IRFs of a battery of survey-based indicators. Expectations of access to credit deteriorate on impact and remain pessimistic for nine months. Employment expectations fall significantly on impact, and expectations for economic growth turn neg-

⁵⁵Recent evidence for the U.S. based on listing-price data shows that list prices can decline within just a couple of weeks after a contractionary monetary policy shock (Gorea, Kryvtsov and Kudlyak, 2023), suggesting that listed prices may adjust even faster than transaction prices.

Figure 10: Monthly IRFs of credit conditions to a monetary policy shock



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. Sample: August 2016–October 2023. Monetary policy shock SD: 4.2bp.

ative almost immediately. Consumer, manufacturing, and services confidence all decline sharply within the first month. Together, these responses point to an immediate and broad-based reassessment: households and firms revise downward their expectations of future income, employment and economic activity well before their contraction actually materialises in the data, regardless of their position in credit markets.⁵⁶

Moreover, we should expect households that are more sensitive to monetary policy and financial news to respond more. To dig into this, we exploit cross-sectional variation in mortgage status within our BBVA sample⁵⁷—households with a mortgage having direct, immediate exposure to the rising borrowing costs and tightening credit documented above, making them a natural, literature-motivated group whose spending is sensitive to the signals documented above.⁵⁸ We estimate local projections at the customer-month level, interacting the monetary policy shock with an indicator for mortgage status and controlling for customer fixed effects.⁵⁹


⁵⁶In Appendix F.3 we further document that survey respondents expect higher unemployment and lower income and house prices. Expectations are comparable across survey respondents in different sectors of the economy.

⁵⁷Recall that details on the construction of the monthly customer-level panel are provided in Section 2.

⁵⁸Other dimensions of cross-sectional heterogeneity can be expected to be relevant in the transmission of monetary policy, but it is beyond the scope of this paper to assess them. For instance, households with workers employed on temporary contracts may respond strongly at short lags through an expectations channel: as shown in Figure 11, employment expectations deteriorate on impact, and temporary workers face a higher near-term probability of dismissal consistent with those expectations.

⁵⁹Appendix F.1 provides full details on the empirical specification and sample restrictions. As a caveat, a

Figure 11: Monthly IRFs of expectations and confidence to a monetary policy shock



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. For expectations for economic growth and harder access to credit the sample starts in April 2020; all other series start in August 2016. All series end in October 2023. Monetary policy shock SD: 4.2bp for full-sample series, 3.6bp for the two shorter series.

Table 2 reports the differential consumption response of mortgagors relative to non-mortgagors from a specification that includes both customer and time fixed effects—so that the aggregate shock effect is fully absorbed and only the differential response is identified. Mortgagors cut spending by an additional 0.12 percentage points relative to non-mortgagors, the peak of the differential, at the 3 month horizon, consistent with the findings of [Cloyne, Ferreira and Surico \(2020\)](#) for the UK and the US.

Note, however, that the fast drop in consumption is not driven by mortgagors alone. As shown in Figure F1 in the appendix, non-mortgagors also reduce consumption significantly at short lags. This suggests that channels beyond direct credit exposure play a complementary role—plausibly reflecting the broad-based deterioration in employment and macroeconomic expectations documented in Figure 11 above.

formal causal analysis of the mortgage channel would require controlling for household leverage, liquidity constraints, and other balance-sheet dimensions—possibly across accounts in different banks and financial institutions—that we do not have information about.

Table 2: Differential Consumption Response: Mortgagors vs. Non-Mortgagors

	Horizon (months)						
	0	1	2	3	6	9	11
Difference	-0.060* (0.053)	-0.043 (0.045)	-0.043* (0.034)	-0.122** (0.037)	-0.024 (0.033)	-0.035** (0.015)	-0.021 (0.024)
Observations (000s)	72130	72130	72130	72130	72130	66957	63516
Customers (000s)	1725	1725	1725	1725	1725	1722	1718
Customer FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Controls: 3 lag(s) consumption</i>							

Notes: LP estimates of differential consumption response (mortgagors minus non-mortgagors) to a one standard deviation monetary policy shock. Dependent variable: 12-month log change in consumption (p.p.). Time FE absorb main shock effect; only the interaction is identified. Driscoll-Kraay SEs. *, **: 68% and 90% CI. Sample: July 2021 – May 2025. The monetary policy shock standard deviation is 2.9bp.

To sum up, the evidence in this section suggests that the fast household spending responses reflect multiple, complementary channels operating at very short lags. Interest rate news is disseminated almost immediately through the media, reaching households within days of the shock. Credit conditions—borrowing costs, credit volumes, asset prices—deteriorate rapidly, directly affecting household balance sheets. And households revise downward their expectations of employment, income, and broader economic activity almost on impact. Consumption falls fast among both mortgagors and non-mortgagors, with the larger decline among mortgagors consistent with their direct exposure to credit market conditions. While we do not advance a claim about which of these channels dominates, our analysis suggests that, together, all these channels contribute to paint a coherent picture of fast monetary transmission to household spending.

7.2 Reconciling Fast Demand Adjustments with Slow Production Adjustments: Inventories

In the previous sections, we have documented that, against a sharp decline in corporate sales within days of a contractionary monetary policy shock, the employment response is initially contained and deepens only over several months. Given that employment adjusts only gradually, and capital is arguably fixed in the short run, production may not fall as fast as sales (even accounting for variations in hours and capacity). By the accounting identity $\text{Output} = \text{Sales} + \Delta \text{Inventories}$, any difference must then show up as inventory accumulation: unsold goods accumulate as demand falls faster than produc-

tion. Over time, however, as firms cut production in line with the weaker demand, they work off excess inventories and inventory growth turns negative. This buffering mechanism is well established in the literature (Ramey and West, 1999): inventories absorb the gap between sales and production when adjustment costs prevent a fast reoptimization of employment, hours and capacity utilization. Empirical evidence is provided by Blander and Maccini (1991), Bils and Kahn (2000), and Wen (2011), as well as by recent work including Rossi (2025), Ferrari (2022), and Marzolf (2024)—the latter contribution shows that inventories actually amplify demand shocks as they propagate upstream through supply chains.


Figure 12 presents three complementary pieces of evidence showing that the inventory buffering mechanism is active in response to monetary policy shocks. The top panel, based on the INE’s Short-term Stock and Inventory Survey, displays the response of aggregate inventory growth in the wholesale and retail sector: positive values indicate accumulation, negative values drawdown, following a monetary policy shock.⁶⁰ Inventories increase on impact—approximately 0.2%—consistent with inventories accumulating as sales fall faster than production adjusts. It stays positive for about three months before turning negative, and then declines steadily, without an interior trough, reaching roughly -2% by the end of the one-year window as firms cut production in line with weaker demand and the employment response deepens (see Figure 2).

The middle and bottom panels in Figure 12 provide complementary evidence from European Commission business surveys (see Section 2 for details) that sharpens the timing of inventory accumulation across the supply chain.⁶¹ The balance of retail firms reporting inventories above versus below normal—shown in the middle panel—is essentially zero on impact but rises within the first month to about 1.5 percentage points and remains elevated, around 1.3 percentage points on average, throughout the year, indicating that managers in downstream firms perceive excess stocks soon after the shock. The bottom panel shows the same balance for manufacturing firms. Unlike the retail response, the manufacturing balance is modestly negative for the first two months before rising sharply to a peak of about 1.4 percentage points around the third month, and remains positive through most of the year. This downstream-first, upstream-later sequencing mirrors the sectoral propagation patterns documented in Section 5.2: retail-

⁶⁰The Encuesta Coyuntural sobre Stocks y Existencias (Base 2021) covers total trade (NACE Section G), and is available from January 2013. The series enters the local projection in log-differences.

⁶¹The European Commission’s harmonised Business and Consumer Surveys for Spain has already been described in Section 2.2 in relation to confidence indicators. The middle panel reports the balance from Question Q2 of the Retail Trade Survey (“Volume of stock currently held”); the bottom panel reports the balance from Question Q4 of the Industry Survey (“Assessment of stocks of finished products”). Both balances are seasonally adjusted and available from January 2000.

Figure 12: Monthly IRFs of Inventories



Notes: LP impulse response functions to a one standard deviation contractionary monetary policy shock. Top panel: growth rate of inventories in wholesale and retail trade (INE). Middle panel: balance of retail trade firms reporting inventories above versus below normal (European Commission). Bottom panel: balance of manufacturing firms reporting stocks of finished products above versus below normal (European Commission). The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample starts in August 2015 and ends in October 2023. The monetary policy shock standard deviation is 3.7bp.

ers accumulate inventories first because consumer demand falls on impact while shipments from manufacturers continue; manufacturers accumulate stocks of finished products later, once retailers cut their orders. The timing is internally consistent with the employment response: employment is essentially flat in the first three months and only begins to fall as manufacturers report excess inventories and cut production.

This evidence suggests that inventories buffer the transition between fast-moving demand and slow-moving production. Our daily corporate sales series captures rapid demand-side adjustments—the immediate response of consumers and firms to tightening financial conditions. Production, proxied by employment, adjusts with a delay, accounting for the time required to change hours and reduce headcounts. The inventory

behavior lends further empirical support to our main claim—that long lags in employment and inflation reflect the slow transmission of fast demand contraction through production adjustment and supply-chain propagation, not a slow initial response of demand to monetary impulses.

7.3 The long lags in inflation match employment

Data on CPI inflation are only available at monthly frequency. Based on data at this frequency, we now show that the response of inflation is smoother and more persistent relative to that of final demand—drawing a pattern similar to that of employment. In the first row of Figure 13 we plot the response of the CPI (right panel) together with the response of employment at monthly frequency, obtained by aggregating our daily data (left panel). Note that the response of employment is very close to our baseline in Figure 2, in line with our results on time aggregation above.


The takeaway from the figure is that, similarly to employment, the CPI decline is gradual and deepens after month 6. Observe that the CPI drops significantly on impact—there is no “price puzzle” in our results. Yet, its short-run response is economically small and slow to set in, relative to the global trough of -0.2 percentage points taking place 11 months after a monetary policy shock.

Furthermore, at the monthly frequency, we can use disaggregated CPI inflation across eleven COICOP categories. As shown in Appendix F.2, for 9 of these categories we find that prices appear not to respond significantly, if at all, at short lags. Instead, broader-based declines set in the second part of the year.⁶² Disaggregated patterns then cumulate up to a smooth and protracted decline in the aggregate CPI, capturing the overall effects of monetary policy on the target measure of inflation. This substantiates our earlier conclusion: while headline macro (demand) aggregates do respond fast, the effects of monetary policy on key variables of interest to central bank—employment and inflation—do materialize at longer lags, in line with the consensus view.

The availability of disaggregated series on employment at monthly frequency also allows us to elaborate further on the transmission of monetary policy. The publicly available series of the Spanish Social Security data distinguish between temporary (fixed-term) contracts, and permanent (no fixed-term declared) contracts. The responses to

⁶²In detail, in Appendix F.2 we report that the short-run CPI responses are driven mainly by contractions in the CPI associated to Transport and Housing and Utilities. Food and non-alcoholic beverages, Alcoholic beverages and tobacco, Furnishings, Equipment and maintenance, and Restaurants and hotels only respond significantly later in the year, 6 months after the monetary policy shock. Finally, we find that even at the one year horizon the response for the CPI of Health, Communication and Education Services remains insignificant.

Figure 13: Monthly IRFs of employment and CPI



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample starts in August 2015 and ends in October 2023. The monetary policy shock standard deviation is 3.7bp.

a policy shock of these two series, shown in Figure 13, are markedly different. The response of permanent employment, which accounts for about 84% of contracts in Spain, is muted at short lags—it becomes significant in statistical and economic terms only in the last months of the year. In contrast, temporary employment adjusts on impact, with a contraction that lasts for around five to six months after the shock. One may observe that the upfront contraction in temporary work is consistent with the pronounced upfront fall in downstream, final consumption sales, as retail and wholesale sector to rely extensively on temporary workers.

The response in total employment shown in the same figure results from the combined dynamics of these two series: the smooth contraction at longer lags is dominated by the loss of jobs with permanent contracts, which account for the largest share of employment; the smaller contraction at short lags results from the drop in temporary employment. Our evidence thus suggests that, in economies with a higher share of temporary/flexible contracts, the response of total employment at short and long lags may well reflect significant composition effects.

Overall, the sluggish response of CPI and employment (in particular that of permanent employment) suggests the relevance of adjustment frictions in both the labor markets and firms' price setting, slowing the transmission of an otherwise fast final demand response—already non-negligible at short lags—to aggregate activity and inflation, only substantial at longer lags.

7.4 Robustness

We close this section by reporting the results from four robustness exercises to which we subject our baseline findings, varying in turn the data sources, the treatment of seasonality and noise, the measurement of the monetary policy shock, and the sample period. Across all four exercises the baseline results are broadly preserved; Appendix G reports each exercise in full.

In our first exercise, we revisit our headline results using publicly available series at monthly frequency. Monthly industrial production—the most common proxy for real activity at this frequency—contracts by 0.30% on impact and by 0.72% at the one-month horizon; the contraction then moderates before becoming statistically significant again through month seven; [Miranda-Agrippino and Ricco \(2021\)](#) and [Bauer and Swanson \(2023\)](#) report comparable short-run responses for the United States—a result suggesting that the pattern is not specific to Spanish data. The VAT-based monthly indicators of final demand also largely mirror the daily baseline: consumption falls in the first month, reaches a trough in the second, and declines again around the tenth month, with consumption services bearing the brunt of the contraction. Investment also drops right away and reaches its low point in the second month, with these short lags in monthly-VAT derived investment appearing less noisy at the monthly than at the daily-transaction level. This pattern is mainly driven by equipment, maintenance, and software, whereas construction reacts more mildly.

Second, we vary the treatment of seasonality and noise in the daily series. The responses of sales, consumption, and employment are similar whether we use our baseline—day-of-the-week and day-of-the-month adjustment combined with Smooth Local Projections—a 30-day backward moving average paired with standard projections, or a model-based TBATS adjustment; only investment is somewhat sensitive to the denoising approach, as for this series achieving a clean separation of the irregular component from noise is challenging. A systematic grid of seasonal-adjustment and smoothing combinations, reported in the appendix, points to the same reading.

Third, we replace the baseline shocks of [Jarociński and Karadi \(2020\)](#) with two al-

ternatives, both drawn from the high-frequency database of [Altavilla et al. \(2019\)](#): the observed one-month OIS change, to which we apply the poor man’s exclusion to control for the information effect, and the Policy Target factor. The implied responses stay within the confidence bands of the baseline, and treating the two series as instruments rather than shocks in a local-projection IV delivers similar estimates.

Fourth, we ask whether the COVID-19 episode drives disproportionately our results. Re-estimated on the daily sample ending in December 2019—which leaves 24 shocks and horizons of up to 93 days—the responses of consumption and employment remain qualitatively similar to the baseline, only slightly attenuated, and the disaggregated COICOP categories stay aligned; a much longer monthly sample, reaching back as far as 2000 for most variables, is likewise consistent.

Taken together, these exercises indicate that our findings on the transmission of monetary policy to daily real activity are broadly stable across data sources, methods of seasonal adjustment and smoothing, measures of the policy shock, and sample periods, and do not hinge on the inclusion of the pandemic in the sample.

8 Conclusion

How fast—and in what sequence—do different parts of the economy respond to monetary policy shocks? To answer this question, we assemble a dataset that observes economic activity at the same frequency at which monetary policy shocks are identified around policy announcements, and provide evidence that such data is necessary to provide a reliable answer. We summarize our findings in three points.

First, demand adjusts fast, and production follows more gradually. Consumption and sales contract within weeks of a monetary policy shock, whereas employment and prices respond only with a delay; investment also declines, albeit more noisily, and its decline is persistent over time. The long lags in employment and inflation commonly associated with Friedman’s dictum therefore reflect the slow transmission of an initially rapid contraction in demand through the production process, rather than a sluggish response of demand itself.

Second, the transmission of monetary impulses unfolds in three stages. Households learn about the shock almost immediately: media coverage of interest rate news increases sharply and its tone turns more negative within days, asset prices adjust contemporaneously, and expectations deteriorate broadly within the first month (Section 7.1). Spending then adjusts where economic theory would predict, falling by more for durable, semi-durable and luxury goods than for necessities, and by more for mortgagors than

for non-mortgagors, though non-mortgagors also cut expenditure at short lags, so that credit exposure is an important channel but not the only one (Sections 5.1 and 7.1). The contraction in final demand then propagates only gradually, from downstream, final-demand sectors to upstream producers, through inventory adjustment and input-output linkages (Sections 7.2 and 5.2). Only at these longer lags does the contraction materialize in aggregate employment and prices.

Third, temporal aggregation is a key reason why these short lags were not detected in earlier work relying on quarterly data. Aggregating our daily series to the quarterly frequency shifts the estimated responses to longer horizons: effects that are significant within the first quarter disappear, and re-emerge only in the following quarter. Our Monte Carlo analysis shows that this distortion is structural rather than a small-sample artifact (Section 6). While the econometric literature has long recognized that temporal aggregation distorts estimated dynamics, our contribution is to show that it can alter a substantive economic conclusion—that demand responds slowly to monetary policy—which has in turn shaped how the profession models transmission.⁶³

Taken together, these findings have implications for how monetary transmission is modeled. DSGE models are routinely calibrated to replicate a delayed demand response estimated using quarterly data, typically through habit formation or adjustment costs; a delay imposed by assumption may misrepresent the sequencing we document. Our results instead call for models that account for the rapid adjustment of demand alongside the gradual transmission of that adjustment to employment, upstream production and prices. Progress in this direction should be facilitated by the increasing availability of high-frequency data. Extending and replicating our analysis across countries and institutional settings would contribute to a fuller account of this transmission, while transaction-level and administrative microdata make it possible to develop structural models that operate at the frequency of policy decisions—models that could offer a timely and granular assessment of how the effects of monetary policy evolve between successive policy meetings.

⁶³The same concern applies wherever adjustment occurs at a higher frequency than the available data, including recent studies that identify transmission mechanisms from annual administrative records.

References

- Agencia Tributaria.** 2023. “Metodologia da la estadistica de ventas diarias del sistema de suministro inmediato de información (SII).” Ministerio de hacienda y función publica, Agencia Tributaria.
- Almgren, Mattias, José-Elías Gallegos, John Kramer, and Ricardo Lima.** 2022. “Monetary policy and liquidity constraints: Evidence from the euro area.” *American Economic Journal: Macroeconomics*, 14(4): 309–340.
- Altavilla, Carlo, Luca Brugnolini, Refet S Gürkaynak, Roberto Motto, and Giuseppe Ragusa.** 2019. “Measuring euro area monetary policy.” *Journal of Monetary Economics*, 108: 162–179.
- Altavilla, Carlo, Refet S. Gürkaynak, Thilo Kind, and Luc Laeven.** 2025. “Monetary transmission with frequent policy events.” European Central Bank ECB Working Paper Series 3157.
- Amemiya, Takeshi, and Roland Y. Wu.** 1972. “The effect of aggregation on prediction in the autoregressive model.” *Journal of the American Statistical Association*, 67(339): 628–632.
- Andersen, Asger Lau, Amalie Sofie Jensen, Niels Johannesen, Claus Thustrup Kreiner, Søren Leth-Petersen, and Adam Sheridan.** 2021. “How do households respond to job loss? Lessons from multiple high-frequency data sets.” CEPR Press Discussion Paper No. 16131.
- Andersen, Asger Lau, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan.** 2022. “Consumer responses to the COVID-19 crisis: Evidence from bank account transaction data.” *The Scandinavian Journal of Economics*, 124(4): 905–929.
- Antràs, Pol, Davin Chor, Thibault Fally, and Russell Hillberry.** 2012. “Measuring the upstreamness of production and trade flows.” *American Economic Review*, 102(3): 41216.
- Ashwin, Julian, Eleni Kalamara, and Lorena Saiz.** 2024. “Nowcasting Euro Area GDP with News Sentiment: A Tale of Two Crises.” *Journal of Applied Econometrics*, 39(5): 887–905.
- Barnichon, Regis, and Christian Brownlees.** 2019. “Impulse response estimation by smooth local projections.” *Review of Economics and Statistics*, 101(3): 522–530.

- Barnichon, Regis, and Geert Mesters.** 2024. "The Phillips curve multiplier." *Journal of Monetary Economics*, 142: 1–18.
- Basu, Susanto.** 1995. "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare." *The American Economic Review*, 85(3): 512–531.
- Bauer, Michael D., and Eric T. Swanson.** 2023. "A reassessment of monetary policy surprises and high-Frequency identification." *NBER Macroeconomics Annual*, 37(1): 87–155.
- Baumeister, Christiane, Danilo Leiva-León, and Eric Sims.** 2021. "Tracking weekly state-level economic conditions." *The Review of Economics and Statistics*, 1–45.
- Bernanke, Ben S, and Kenneth N Kuttner.** 2005. "What explains the stock market's reaction to Federal Reserve policy?" *The Journal of finance*, 60(3): 1221–1257.
- Bernanke, Ben S., and Mark Gertler.** 1995. "Inside the Black Box: The Credit Channel of Monetary Policy Transmission." *Journal of Economic Perspectives*, 9(4): 27–48.
- Bernanke, Ben S, Jean Boivin, and Piotr Elias.** 2005. "Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach." *The Quarterly journal of economics*, 120(1): 387–422.
- Bils, Mark, and James A. Kahn.** 2000. "What Inventory Behavior Tells Us about Business Cycles." *American Economic Review*, 90(3): 458–481.
- Blinder, Alan S., and Louis J. Maccini.** 1991. "Taking Stock: A Critical Assessment of the Microeconomic Inventory Literature." *Journal of Economic Perspectives*, 5(1): 73–96.
- Bounie, David, Youssouf Camara, Etienne Fize, John Galbraith, Camille Landais, Chloe Lavest, Tatiana Pazem, Baptiste Savatier, et al.** 2020. "Consumption dynamics in the covid crisis: real time insights from French transaction bank data." *Covid Economics*, 59: 1–39.
- Buda, Gergely, Stephen Hansen, T Rodrigo, VM Carvalho, Á Ortiz, and JVR Mora.** 2022. "National accounts in a World of naturally occurring data: A proof of concept for consumption." CEPR Press Discussion Paper No. 17519.
- Burr, Natalie, and Tim Willems.** 2024. "About a rate of (general) interest: how monetary policy transmits." Quarterly Bulletin, Bank of England.

- Calza, Alessandro, Tommaso Monacelli, and Livio Stracca.** 2013. "Housing Finance and Monetary Policy." *Journal of the European Economic Association*, 11(s1): 101–122.
- Carroll, Christopher D.** 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics*, 118(1): 269–298.
- Carvalho, Carlos.** 2006. "Heterogeneity in price stickiness and the real effects of monetary shocks." *B.E. Journal of Macroeconomics*, 2(1): 1–58.
- Chetty, Raj, John N Friedman, Nathaniel Hendren, Michael Stepner, et al.** 2020. "The economic impacts of COVID-19: Evidence from a new public database built using private sector data." NBER Working Paper No. 27431.
- Christiano, Lawrence J., and Martin Eichenbaum.** 1987. "Temporal aggregation and structural inference in macroeconomics." *Carnegie-Rochester Conference Series on Public Policy*, 26: 63–130.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans.** 1999. "Chapter 2 Monetary policy shocks: What have we learned and to what end?" In . Vol. 1 of *Handbook of Macroeconomics*, 65–148. Elsevier.
- Cieslak, Anna, and Andreas Schrimpf.** 2019. "Non-monetary news in central bank communication." *Journal of International Economics*, 118: 293–315.
- Cloyne, James, Clodomiro Ferreira, and Paolo Surico.** 2020. "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism." *The Review of Economic Studies*, 87(1): 102–129.
- Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber.** 2022. "Monetary Policy Communications and Their Effects on Household Inflation Expectations." *Journal of Political Economy*, 130(6): 1537–1584.
- Corsetti, Giancarlo, Joao B Duarte, and Samuel Mann.** 2022. "One money, many markets." *Journal of the European Economic Association*, 20(1): 513–548.
- Cuevas, Ángel, Ramiro Ledo, and Enrique M Quilis.** 2021. "Seasonal adjustment of the Spanish sales daily data." *SERIEs*, 12(4): 687–708.
- De Livera, Alysha M, Rob J Hyndman, and Ralph D Snyder.** 2011. "Forecasting time series with complex seasonal patterns using exponential smoothing." *Journal of the American statistical association*, 106(496): 1513–1527.

- Diebold, Francis X.** 2022. "Real-Time Real Economic Activity: Entering and Exiting the Pandemic Recession of 2020." In *Essays in Honour of Fabio Canova*. Vol. 44 of *Advances in Econometrics*, 5–24. Emerald Group Publishing Limited.
- Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao.** 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review*, 107(11): 3550–3588.
- Eraslan, Sercan, and Thomas Götz.** 2021. "An unconventional weekly economic activity index for Germany." *Economics Letters*, 204: 109881.
- Erceg, Christopher, and Andrew Levin.** 2006. "Optimal monetary policy with durable consumption goods." *Journal of monetary Economics*, 53(7): 1341–1359.
- Ferrari, Alessandro.** 2022. "Inventories, Demand Shocks Propagation and Amplification in Supply Chains." *arXiv*. Shows how inventories amplify demand shocks upstream and slow production adjustment.
- Ferreira, Clodomiro, José Miguel Leiva, Galo Nuño, Álvaro Ortiz, Tomasa Rodrigo, and Sirenia Vazquez.** 2022. "The heterogeneous impact of inflation on households' balance sheets." Bank of Spain Working Paper No. 2403.
- Fieldhouse, Andrew, Karel Mertens, and Morten O. Ravn.** 2018. "The Macroeconomic Effects of Government Asset Purchases: Evidence from Postwar U.S. Housing Credit Policy." *The Quarterly Journal of Economics*, 133(3): 1503–1560.
- Flodén, Martin, Matilda Kilström, Jósef Sigurdsson, and Roine Vestman.** 2021. "Household Debt and Monetary Policy: Revealing the Cash-Flow Channel." *Economic Journal*, 131(636): 1742–1771.
- Friedman, Milton.** 1961. "The lag in effect of monetary policy." *Journal of Political Economy*, 69(5): 447–466.
- Ganong, Peter, and Pascal Noel.** 2019. "Consumer spending during unemployment: Positive and normative implications." *American economic review*, 109(7): 2383–2424.
- Gertler, Mark, and Peter Karadi.** 2015. "Monetary policy surprises, credit costs, and economic activity." *American Economic Journal: Macroeconomics*, 7(1): 44–76.
- Geweke, John.** 1978. "Temporal aggregation in the multiple regression model." *Econometrica: Journal of the Econometric Society*, 643–661.

- Ghassibe, Mishel.** 2021. "Monetary policy and production networks: an empirical investigation." *Journal of Monetary Economics*, 119: 21–39.
- Gorea, Denis, Oleksiy Kryvtsov, and Marianna Kudlyak.** 2023. "House Prices Respond Promptly to Monetary Policy Surprises." *FRBSF Economic Letter*, , (2023-03).
- Grigoli, Francesco, and Damiano Sandri.** 2026. "Monetary policy and credit card spending." *European Economic Review*, 185: 105294.
- Gürkaynak, RS, B Sack, and ET Swanson.** 2005. "Do actions speak louder than words? The response of asset prices to monetary policy actions and statements." *International Journal of Central Banking*, 1(1): 55–93.
- Hanson, Samuel G, and Jeremy C Stein.** 2015. "Monetary policy and long-term real rates." *Journal of Financial Economics*, 115(3): 429–448.
- Harvey, Andrew C.** 1989. *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press.
- Holm, Martin Blomhoff, Pascal Paul, and Andreas Tischbirek.** 2021. "The transmission of monetary policy under the microscope." *Journal of Political Economy*, 129(10): 2861–2904.
- Ishan, Anand, Valerie A. Ramey, and Peter J. Klenow.** 2024. "The Aggregate Effects of Global Warming." National Bureau of Economic Research NBER Working Paper 32164.
- Jacobson, Margaret, Christian Matthes, and Todd Walker.** 2026. "Temporal Aggregation Bias and Monetary Policy Transmission." *Journal of Political Economy Macroeconomics*, forthcoming.
- Jarociński, Marek.** 2024. "Estimating the Feds unconventional policy shocks." *Journal of Monetary Economics*, 144: 103548.
- Jarociński, Marek, and Peter Karadi.** 2020. "Deconstructing monetary policy surprises—the role of information shocks." *American Economic Journal: Macroeconomics*, 12(2): 1–43.
- Jordà, Òscar.** 2005. "Estimation and inference of impulse responses by local projections." *American economic review*, 95(1): 161–182.
- Jordà, Òscar, and Alan M. Taylor.** 2025. "Local Projections." *Journal of Economic Literature*, 63(1): 59–110.

- Kuttner, Kenneth N.** 2001. "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market." *Journal of monetary economics*, 47(3): 523–544.
- Lamla, Michael J., and Dmitri V. Vinogradov.** 2019. "Central Bank Announcements: Big News for Little People?" *Journal of Monetary Economics*, 108: 21–38.
- Lane, Philip R.** 2024. "The analytics of the monetary policy tightening cycle." Guest lecture by Philip R. Lane, Member of the Executive Board of the ECB, at Stanford Graduate School of Business.
- Leetaru, Kalev, and Philip A. Schrodtt.** 2013. "GDELT: Global Data on Events, Location, and Tone, 1979–2012." Vol. 2.
- Lenza, Michele, and Giorgio E Primiceri.** 2022. "How to estimate a vector autoregression after March 2020." *Journal of Applied Econometrics*, 37(4): 688–699.
- Lewis, Daniel J.** 2023. "Announcement-specific decompositions of unconventional monetary policy shocks and their effects." *Review of Economics and Statistics*, 1–46.
- Lewis, Daniel J., Christos Makridis, and Karel Mertens.** 2019. "Do monetary policy announcements shift household expectations?" Federal Reserve Bank of Dallas Working Papers 1906.
- Lewis, Daniel J, Karel Mertens, James H Stock, and Mihir Trivedi.** 2022. "Measuring real activity using a weekly economic index." *Journal of Applied Econometrics*, 37(4): 667–687.
- Mackowiak, Bartosz, and Mirko Wiederholt.** 2009. "Optimal sticky prices under rational inattention." *American Economic Review*, 99(3): 769–803.
- Marcet, Albert.** 1991. "Temporal aggregation of economic time series." In *Rational Expectations Econometrics*. 237–282. CRC Press.
- Marzolf, M. J.** 2024. "Retail & Wholesale Inventories: A Literature Review and Path Forward." *Journal of Business Logistics (JBL)*. Retailers with higher inventory turnover respond faster to demand shocks, highlighting buffer role.
- McKay, Alisdair, and Johannes F Wieland.** 2021. "Lumpy durable consumption demand and the limited ammunition of monetary policy." *Econometrica*, 89(6): 2717–2749.
- Miranda-Agrippino, Silvia, and Giovanni Ricco.** 2021. "The transmission of monetary policy shocks." *American Economic Journal: Macroeconomics*, 13(3): 74–107.

- Monacelli, Tommaso.** 2009. "New Keynesian models, durable goods, and collateral constraints." *Journal of Monetary Economics*, 56(2): 242–254.
- Montiel Olea, José Luis, and Mikkel Plagborg-Møller.** 2021. "Local projection inference is simpler and more robust than you think." *Econometrica*, 89(4): 1789–1823.
- Mountford, Andrew, and Harald Uhlig.** 2009. "What are the effects of fiscal policy shocks?" *Journal of Applied Econometrics*, 24(6): 960–992.
- Nakamura, Emi, and Jón Steinsson.** 2018. "High-frequency identification of monetary non-neutrality: the information effect." *The Quarterly Journal of Economics*, 133(3): 1283–1330.
- Nakamura, Emi, and Jón Steinsson.** 2010. "Monetary non-neutrality in a multisector menu cost model." *The Quarterly Journal of Economics*, 125(3): 961–1013.
- Ozdagli, Ali, and Michael Weber.** 2017. "Monetary policy through production networks: Evidence from the stock market." National Bureau of Economic Research.
- Pasten, Ernesto, Raphael Schoenle, and Michael Weber.** 2020. "The propagation of monetary policy shocks in a heterogeneous production economy." *Journal of Monetary Economics*, 116: 1–22.
- Ramey, Valerie A.** 2016. "Macroeconomic shocks and their propagation." *Handbook of macroeconomics*, 2: 71–162.
- Ramey, Valerie A., and Kenneth D. West.** 1999. "Inventories." In *Handbook of Macroeconomics*. Vol. 1B, , ed. John B. Taylor and Michael Woodford, 863–923. Elsevier.
- Romer, Christina D, and David H Romer.** 2004. "A new measure of monetary shocks: Derivation and implications." *American Economic Review*, 94(4): 1055–1084.
- Rossi, Luca.** 2025. "Uncovering the InventoryBusiness Cycle Nexus." Bank of Italy Temi di Discussione. Finds that sales forecast errors drive short-run inventory investment (buffer-stock evidence).
- Rua, António, and Nuno Lourenço.** 2020. "The DEI: Tracking economic activity daily during the lockdown." Banco de Portugal, Working Papers 2020-13.
- Schorfheide, Frank, and Dongho Song.** 2024. "Real-Time Forecasting with a (Standard) Mixed-Frequency VAR During a Pandemic." *International Journal of Central Banking*, 20(4): 275–320.

- Sims, Christopher A.** 1971. "Discrete approximations to continuous time distributed lags in econometrics." *Econometrica*, 39(3): 545–563.
- Slacalek, Jiri, Oreste Tristani, and Giovanni L Violante.** 2020. "Household balance sheet channels of monetary policy: A back of the envelope calculation for the euro area." *Journal of Economic Dynamics and Control*, 115: 103879.
- Sterk, Vincent, and Silvana Tenreyro.** 2018. "The transmission of monetary policy through redistributions and durable purchases." *Journal of Monetary Economics*, 99: 124–137.
- Stock, James H.** 1987. "Temporal aggregation and structural inference in macroeconomics a comment." Vol. 26, 131–139, North-Holland.
- Stock, James H, and Mark W Watson.** 2018. "Identification and estimation of dynamic causal effects in macroeconomics using external instruments." *The Economic Journal*, 128(610): 917–948.
- Swanson, Eric T.** 2021. "Measuring the effects of federal reserve forward guidance and asset purchases on financial markets." *Journal of Monetary Economics*, 118: 32–53.
- Wen, Yi.** 2011. "Input and Output Inventory Dynamics." *American Economic Journal: Macroeconomics*, 3(4): 181–212.
- Woodford, Michael.** 2003. *Interest and prices*. Princeton, NJ [u.a.]:Princeton Univ. Press.

Online Appendix

A Appendix for Section 2

A.1 Summary of Data Collection and Sources

Table A1 presents an overview of the data used in the main body of the paper, omitting disaggregation of the different series of sales and consumption by subcategories. Note that while different series have different starting dates, all end in October 2023. For the prices block of the dataset, we only have monthly data; for other variables, we have both daily and monthly data.

Variable	Proxy	Source	Frequency	Start date
Real activity				
Sales & output	Sales	Spanish Tax Authority	Daily / Monthly	July 1st, 2017 / January 2000
	IP	INE	Monthly	January 2000
Consumption	Private consumption	BBVA	Daily	August 1st, 2015
	Private consumption	Spanish Tax Authority	Monthly	January 2000
Investment	Investment	BBVA	Daily	April 6th, 2017
	Investment	Spanish Tax Authority	Monthly	January 2000
Employment	Employment	Spanish Social Security	Daily	August 3rd, 2015
Inventories	Stock index, total trade	INE	Monthly	January 2013
	Survey balance, retail (Q2)	EU Commission	Monthly	January 2000
	Survey balance, manufacturing (Q4)	EU Commission	Monthly	January 2000
Financial Markets				
Interest rate	Euribor	European Money Markets Institute	Daily	January 4th, 1999
	Interest rates for housing	Bank of Spain (Statistics Bulletin)	Monthly	January 2003
Stock prices	IBEX35	Bloomberg	Daily	January 3rd, 2005
Prices				
Consumer prices	CPI	INE	Monthly	January 2000
Housing prices	Average price per square meter	CIEN	Monthly	January 2007
Expectations				
Inflation expectations	Inflation-linked swaps	Bloomberg	Daily	June 3rd, 2004
Real activity expectations	Consumer sentiment indicators	EU Commission	Monthly	January 2000
	Business sentiment indicators	EU Commission	Monthly	January 2000
	Consumer expectations	ECB	Monthly	April 2020
Financial markets expectations	Consumer expectations	ECB	Monthly	April 2020

Table A1: Data overview

A.2 Descriptive Statistics

Table A2 reports descriptive statistics of year-on-year growth rates of 30-day backward-looking moving averages for real sales, consumption, investment and employment. In our baseline sample, on average total real sales increased 3.09%, total real consumption 2.41%, real investment 6.72% and employment 2.31%. Employment is the least volatile time series, while investment is the most volatile series, in line with the standard ranking of volatilities in lower frequency data. Reflecting the effects of the COVID-19 pandemic, in our baseline sample the series of sales and consumption exhibit large variations, with a maximum contraction as deep as, respectively, -41.7% and -29.5%. Note nonetheless that, in our findings, sales and consumption fall at most by 0.8% and 0.4% in response to a contractionary monetary shock. This suggests that monetary policy shocks played a very minor role in determining the total variation in consumption and sales in our sample, as commonly found in the literature.

Table A2: Descriptive statistics of the main variables

	Mean	SD	Min	Max
Sales	3.09	13.02	-41.68	39.93
Consumption	2.41	7.08	-29.48	27.05
Investment	6.72	18.09	-50.40	53.59
Employment	2.31	2.21	-4.74	5.07

Notes: Sales, consumption, investment and employment are measured as YoY growth rates of their 30-day moving averages. We deflate daily sales and consumption using the monthly Consumer Price Index (CPI); and daily investment using the monthly investment sales price deflator used by the Spanish Tax Authority.

Turning to the disaggregated consumption and sales series, Table A3 lists the COICOP consumption subaggregates that we use in our analysis, while Table A4 reports summary statistics. Table A5 lists the NACE classification of sales data by sector from the Spanish Tax Authority—the corresponding descriptive statistics are presented in Table A6. Note that in our sample, seven sectors—beverages and tobacco, textile, paper, chemical industry, pharmaceutical, rubber and plastics, and electronics—experienced negative mean real growth rates in sales. The weighted average growth of sales by category aggregates to the total sales growth in Table A2.

Table A3: COICOP consumption categories (two-digit)

Category	Description
01	Food and Non-Alcoholic Beverages
02	Alcoholic Beverages, Tobacco, and Narcotics
03	Clothing and Footwear
04	Housing, Water, Electricity, Gas, and Other Fuels
05	Furnishings, Household Equipment, and Routine Household Maintenance
06	Health
07	Transport
08	Communication
09	Recreation and Culture
10	Education
11	Restaurants and Hotels

Notes: This table displays the 11 COICOP categories we use for classifying consumption transactions. In line with the Spanish Statistical Office, we use the European COICOP system in place of the international COICOP system. The main difference is that the latter has two separate categories for insurance and financial services and personal care, social protection and miscellaneous goods and services which in ECOICOP are merged into a single Miscellaneous Goods and Services category.

Table A4: Descriptive statistics, COICOP consumption categories (two-digit)

Two-Digit Category	Mean	SD	Min	Max
01	8.54	12.73	-14.75	51.22
02	2.05	9.72	-30.33	40.80
03	2.69	19.42	-59.48	121.79
04	0.37	9.51	-22.01	22.27
05	4.42	11.42	-32.95	57.58
06	11.43	17.35	-45.51	115.21
07	7.65	28.11	-70.63	204.80
08	0.77	5.69	-10.31	20.80
09	3.48	18.26	-60.60	101.09
10	5.44	20.30	-53.35	115.33
11	7.04	27.84	-71.52	199.21

Notes: Consumption categories are measured as YoY growth rates of their 30-day moving averages. Categories are deflated using the CPI at their corresponding COICOP category.

Table A5: Sales sectors, NACE code, and description

Sector	NACE code	Description
Manufacturing: Food	C10	Manufacture of food products
Manufacturing: Beverages and tobacco	C11 + C12	Manufacture of beverages and tobacco
Manufacturing: Textile	C13 + C14 + C15	Manufacture of textiles, wearing apparel and leather products
Manufacturing: Paper	C17 + C18	Paper industry and graphic arts
Manufacturing: Chemical industry	C20	Chemical industry
Manufacturing: Pharmaceutical	C21	Manufacturing of pharmaceutical products
Manufacturing: Rubber and plastics	C22 + C23	Manufacturing of rubber and plastics and other non-metallic mineral products
Manufacturing: Metallurgy	C24 + C25	Metallurgy and manufacturing of iron, steel, ferroalloys, and metal products (except machinery and equipment)
Manufacturing: Electronics	C26 + C27	Manufacture of computer, electronic, optical products and of electrical equipment
Manufacturing: Motor vehicles, trailers, and semi-trailers	C29	Manufacturing of motor vehicles, trailers, and semi-trailers
Manufacturing: Furniture	C16 + C31	Wood and cork industry; manufacturing of furniture
Manufacturing: Machinery	C28 + C30 + C33	Manufacturing of machinery and equipment; manufacturing of other transport equipment; repair and installation of machinery and equipment
Manufacturing: Coking and oil refining	C19 + C32	Coking and oil refining; other manufacturing industries
Energy	D	Electricity, gas, steam and air conditioning supply
Construction	F	Construction
Wholesale and retail trade	G	Wholesale and retail trade; repair of motor vehicles and motorcycles
Transportation and storage	H	Transportation and storage
Hospitality	I	Lodging, food and beverage services, event planning, theme parks, travel agency, tourism, hotels, restaurants, nightclubs, and bars
Information and communication	J	Information and Communication
Professional, scientific and administrative	M + N	Professional, scientific and technical activities

Table A6: Descriptive statistics of sales by sector

Sector	Mean	SD	Min	Max
Manufacturing: Food	0.41	10.45	-29.15	26.86
Manufacturing: Beverages and tobacco	-0.76	15.67	-46.51	70.95
Manufacturing: Textile	-0.69	25.96	-71.27	193.32
Manufacturing: Paper	-0.73	7.74	-23.02	32.65
Manufacturing: Chemical industry	-2.22	8.49	-23.35	27.99
Manufacturing: Pharmaceutical	-2.42	14.87	-33.04	40.49
Manufacturing: Rubber and plastics	-0.07	15.99	-45.12	86.23
Manufacturing: Metallurgy	0.33	18.73	-51.51	107.46
Manufacturing: Electronics	-2.59	19.55	-58.50	110.49
Manufacturing: Motor vehicles, trailers, and semi-trails	9.38	83.27	-91.97	963.87
Manufacturing: Furniture	2.61	21.90	-53.44	147.43
Manufacturing: Machinery	1.82	18.92	-38.82	59.51
Manufacturing: Coking and oil refining	0.55	26.41	-68.05	100.81
Energy	12.05	42.51	-55.96	145.36
Construction	4.81	17.65	-32.89	60.58
Wholesale and retail trade	3.91	12.23	-36.17	56.41
Transport and storage	5.92	28.93	-48.11	96.33
Hospitality	31.32	76.50	-91.92	300.83
Information and communication	1.42	9.44	-22.20	39.05
Professional and administrative services	3.29	14.40	-46.26	51.75

Notes: Sales by sector measured as YoY growth rates of their 30-day moving averages. Manufacturing and Construction sectors are deflated using the closest producer price indexes (PPI); Wholesale and Retail Trade, and Transport and Storage with the closest disaggregated CPI. Finally, we use the Services Price Index (SPI) to deflate the remaining (service) sectors.

A.3 Additional Data Details: Investment

Our baseline daily series on sales and employment, along with discussions of the underlying data sources and methods, are publicly available from Spanish public institutions – e.g. the national Tax Authority or the Ministry for Social Security. Our daily aggregate consumption, is instead discussed in detail by Buda et al. (2022). Our daily aggregate investment series, instead, has not been previously released or analysed. Hereafter, we complement the discussion in the main text with additional details on the steps taken to construct this series, and compare it to other investment series available at lower frequency from the national accounts or tax declarations.

Recapping the discussion in the main text, we evaluate all transfers and reserve factoring⁶⁴ operations mediated by BBVA. We keep only transactions where both parties can be identified as firms, either because both ends of the transaction are linked to BBVA corporate accounts, or because the BBVA party is a firm and we can identify the counterpart as another firm through name-matching with the SABI dataset—which includes the universe of all Spanish firms. In either case, both firms are always associated with a CNAE sector code. We include only transactions related to a payment for goods or services, as can be typically be inferred from presence of the keyword ‘invoice’ or variations thereof in the concept of the payment. We eliminate all transactions where both parties belong to the same ownership group (as the transaction may reflect transfers of funds and not a purchase). Likewise, we eliminate all transactions where one of the parties belongs to the financial sector, the public sector, or a non-profit foundation.

For each of these firm-to-firm transactions in our dataset, we then assign the probability that the transaction refers to the purchase of an investment good. As explained in the main text, we do so by classifying each firm (either selling or purchasing) in a sector, reweighing our data and exploiting the share of investment spending in each sector derived from the 2019 Input-Output table for Spain, made available by INE.⁶⁵ Below we provide further discussion of how we transform our data to make sure that (a) they are representative of aggregate corporate sector in the Spanish economy and (b) they appropriately capture investment spending

Formally, let Y_i^j stand for the sales of intermediate goods from sector i to j , and I_i for the sales of investment goods by firms in sector i , both being the sum of the respective annual flows during 2019 as recorded in the I-O table. Moreover, let Z_i^j denote the sum of all sales from sector i to firms of sector j recorded in BBVA during 2019 that have passed the selection procedure described above. Finally, let Z_i denote the total corporate sales –

⁶⁴In Spanish banking these are also called "Confirming" operations.

⁶⁵Tabla Origen Destino 2019 from INE. https://www.ine.es/daco/daco42/cne15/cne_tod_19.xlsx

recorded by BBVA – by firms of sector i .

Concerning the first issue above in (a), the sectoral distribution of BBVA corporate clients maybe biased relative to the Spanish sectoral distribution of activity. This will be the case whenever the relative size of the sectors as mediated by BBVA $\left(\frac{Z_i}{\sum_{\forall k} Z_k}\right)$ is different from the IO table $\left(\frac{Y_i+I_i}{\sum_{\forall k} (Y_k+I_k)}\right)$. We correct the potential bias by reweighing our data as follows:

$$\phi_i = \frac{Z}{Y+I} \times \frac{\frac{Z_i}{Z}}{\frac{Y_i+I_i}{Y+I}} \quad (4)$$

The first factor corrects the size of the recorded sales across firms in the BBVA account ($Z = \sum_{\forall k} Z_k$) and expands it to cover all sales in the Spanish economy ($Y+I = \sum_{\forall k} (Y_k+I_k)$). The second factor corrects the relative size of the sector in BBVA with the size of the sector in the Spanish economy. This correction implies that $\frac{Z_i}{\phi_i}$ is now an estimate of what we would expect to be the total sales from sector i to j in Spain, by projecting the amount observed in the BBVA data with appropriate weights.⁶⁶

Concerning the second issue in (b), we need to estimate the probability that BBVA transactions between two firms (in any two sectors) reflects the sale of an investment good. Recall that from official input-output data, we know the total sales of intermediate goods from i to j in a given year, Y_i^j . Thus, letting δ_i^j denote the percentage of sales on investment goods, we can write $(1 - \delta_i^j)Z_i^j \frac{1}{\phi_i} = Y_i^j$ such that:⁶⁷

$$\delta_i^j = 1 - \frac{Y_i^j}{Z_i^j} \phi_i. \quad (5)$$

The investment sales from sector i to j attributed by our algorithm is then simply:

$$\hat{I}_i^j(t) = \delta_i^j \times Z_i^j(t) \quad (6)$$

where $Z_i^j(t)$ is the sum of all sales during day t by firms of sector i to firms of sector

⁶⁶Notice that ϕ_i is not the ideal reweighting, as it does not take into account the distribution of sales across purchasing sectors. That is, if the data were available we would use a more accurate re-weighting factor such as: $\phi_i^j = \frac{Z_i^j}{Y_i^j+I_i^j}$. Unfortunately, that is not possible, as the Input-Output Tables report intermediate goods purchases seller-buyer level (this is, they report Y_i^j), but fail to report the equivalent number for the sales of investment goods (this is, they report I_i , but not I_i^j). We opt to use the best available weight ϕ_i , and maintain the assumption that within each cell, the share of intermediate input to investment is approximately constant. This approximation turns out to be good enough given our validation exercises.

⁶⁷We further check that the implied δ_i^j estimate is a number between 0 and 1 and renormalize the data whenever this is not the case.

j mediated by BBVA. Summing up, we obtain our proxy for aggregate investment on a given day t :

$$\hat{I}(t) = \sum_{\forall i,j} \hat{I}_i^j(t) \quad (7)$$

to which we apply standard filters for outliers.

We build this high-frequency indicator of investment precisely because there are no high-frequency measures from administrative records that we can use for our purposes. We can however validate our series by aggregating our data to lower frequencies and comparing it to two investment series released by Spanish authorities: the quarterly series from the national accounts – released by the National Statistics Institute – and the monthly aggregate investment series compiled by the Spanish tax authority (AEAT), based on tax declarations of corporate sales (from VAT records) for large Spanish firms with sales of more than 6 million annual euros. As discussed in the main text, although the number of firms declaring daily sales for VAT purposes is only around 1% of firms, they cover over 60% of total sales. Note that, similarly to our own procedure above, the latter monthly proxy for aggregate investment also relies on the Spanish Input-Output tables in order to classify corporate sales as investment.

Overall, we find that our series tracks both series reasonably well. Figure A1 plots the YoY growth rate of the monthly aggregation of our series (in black) versus the monthly AEAT series (in red) for the time period when both are available. The correlation of the two series is 0.70. In Figure A2 we compare the time series of our investment proxy aggregated to a quarterly frequency (in black) relative to the official series for private investment in the Spanish national accounts (in red). This figure has three panels. In panel a), we compare the levels (the correlation is 0.95), in panel b) the QoQ growth rates (the correlation is 0.90), and in panel c) the YoY growth rates (the correlation is 0.95). Notice that QoQ growth rates of the raw series may be driven by common seasonal patterns, which could inflate the correlation between series. However this argument is less applicable to the exercise in (c), as year-on-year growth rates naturally difference out quarter-specific effects.

We conclude that our high-frequency series appears to be a reliable proxy indicator of investment.

Figure A1: Monthly aggregation of BBVA Investment Indicator and AEAT measure







Figure A2: Quarterly series of BBVA Investment Indicator, and comparison with National Accounts



(a) Levels



(b) Quarter on Quarter Growth



(c) Year on Year Growth

B Appendix for Section 3

B.1 Monetary Policy Shocks: Descriptive Statistics

Table B1 presents descriptive statistics for the monetary policy shocks (MP) in two samples: the baseline sample and the full sample. The baseline sample starts on August 1, 2016 and ends in October 30, 2023, while the full sample starts in January 1, 2000, and ends in the same day as the baseline sample. The table compares the three different measures of MP shocks described in the text—JK (baseline), Target, and 1-month OIS—providing summary statistics for each.⁶⁸ The descriptive statistics indicate that the distribution of the JK MP shocks in the baseline sample is comparable to that of the full sample. In both samples, the JK MP shocks exhibit a distribution that is marginally skewed to the right. Turning to how the alternative monetary policy shocks compare to the baseline shocks, the Target descriptive statistics are quite similar, the 1-month OIS slightly different (a lower variance). One key reason for the discrepancy is the difference in the number of shocks in the series—24 against 55 in the baseline shock series. The difference follows from applying the poor man’s approach, which removes shocks depending on the sign of the comovement between the change in the stock market price index and the change in interest rate around policy announcements.

Table B2 shows a significant correlation among different measures of European Central Bank (ECB) monetary policy surprises within the baseline sample. The 1-month OIS shows a robust positive correlation with both the JK measure (0.75) and the Target shock (0.51), indicating that directly observed market-based and constructed policy rate surprises are closely linked. Similarly, the correlation between the Target and the JK measure is 0.52. This correlation across the different measures of monetary policy surprises helps explain why our findings are consistent across them—see Figure G3.

Figure B1 plots the time series of monetary policy shocks from Jarociński and Karadi (2020), panel (a) for the baseline sample and panel (b) for the full sample. Fluctuations in monetary policy shocks are less pronounced in the first half of the sample compared to the second half. This is because in the first half of the sample policy rates are at their effective (zero) lower bound—while the second half is characterized by the tightening cycle triggered by the sustained rise in inflation. One notable event is the significant contractionary surprise of approximately 19 basis points on March 12, 2020. These shock patterns observed in our baseline sample, coupled with the onset of the COVID crisis in

⁶⁸As explained in Appendix G.3, the JK (baseline) shock is from Jarociński and Karadi (2020), the other two monetary policy shocks (used in our robustness checks) are constructed based on the publicly accessible Euro Area Monetary Policy Event-Study Database (EA-MPD) created by Altavilla et al. (2019).

March 2020, prompt us to perform sub-sampling robustness checks. These robustness checks are detailed in Section G.4, where we examine whether our main conclusions remain valid when restricting the sample to end in December 2019, before the start of COVID-19 and before the major contractionary shock.

Finally, note that, as demonstrated in panel (b) of Figure B1, the degree of variation and the occurrence of more extreme shocks are comparable in the baseline sample (2016-2023) and in the longer sample.

Table B1: MP descriptive statistics

MP shock	N	Mean	Median	SD	Min	Max
Baseline sample						
JK (baseline)	55	0.90	0.31	3.74	-5.38	18.76
Target	55	0.55	-0.30	3.58	-4.96	21.13
1-month OIS	24	0.76	0.03	2.86	-4.45	10.57
Full sample						
JK (baseline)	224	0.51	0.20	3.29	-8.53	18.76


Notes: The baseline sample starts in August 1st, 2016 (following the year-on-year transformation applied to the dependent variables in our local projections) and ends in October 30th, 2023, while the full sample starts in January 1st, 2002 and ends in the same day as the baseline sample.

Table B2: MP cross-correlation in the Baseline sample


	1-month OIS	Target	JK (baseline)
1-month OIS	1.00	0.51	0.75
Target	0.51	1.00	0.52
JK (baseline)	0.75	0.52	1.00

Figure B1: Jarociński and Karadi (2020) monetary policy shocks

(a) Baseline sample



(b) Full sample



Notes: This figure shows the baseline monetary policy shocks (Jarociński and Karadi, 2020) time series we use in this paper. Panel (a) shows the time series of the baseline monetary policy shock in the baseline sample, while Panel (b) shows the time series of the baseline monetary policy shock in the full sample.

B.2 Recovering Impulse-Response Functions in Levels from Impulse-Response Functions in Year-on-year Growth Rates

Consider a time series y_t in log-levels. Assuming y_t is covariance stationary, the Wold representation of the time series is given by:

$$y_t = \sum_{j=0}^{\infty} \phi_j \varepsilon_{t-j} + \eta_t, \quad (8)$$

with $\phi_0 = 1$ and $\sum_{j=1}^{\infty} \phi_j^2 < \infty$, and where ϕ_j are coefficients, ε_{t-j} are uncorrelated innovations, and η_t is a deterministic component.

An impulse response function is defined as the response of variable y to innovation ε_t at horizon $h = 0, 1, \dots, H$. Given the Wold representation of y_t , we have that

$$IRF_h = \frac{\partial y_{t+h}}{\partial \varepsilon_t} = \phi_j. \quad (9)$$

Assuming the frequency of y_t be daily, we can define the year-on-year (YoY) growth rate as $z_t = y_t - y_{t-365}$. Since y_t is covariance stationary, so is z_t , and its Wold representation is given by:

$$z_t = \sum_{j=0}^{\infty} b_j \varepsilon_{t-j} = \sum_{j=0}^{\infty} \phi_j \varepsilon_{t-j} - \sum_{j=0}^{\infty} \phi_j \varepsilon_{t-365-j}. \quad (10)$$

Hence, the impulse response of the YoY is given by

$$IRF_h^{YoY} = \frac{\partial z_{t+h}}{\partial \varepsilon_t} = b_j. \quad (11)$$

We can recover the impulse response function of the variables in log-levels ϕ_j from YoY impulse response functions b_t . For $0 \leq h < 365$:

$$IRF_h^{YoY} = \frac{\partial z_{t+h}}{\partial \varepsilon_t} = b_j = \phi_j = \frac{\partial y_{t+h}}{\partial \varepsilon_t} = IRF_h. \quad (12)$$

The intuition behind this result is straightforward. Changes in YoY growth rates, $z_t = y_t - y_{t-365}$, that are induced by innovations between time 0 and 364 days ago can only be driven by y_t because y_{t-365} cannot be affected by future innovations. Now from 365 days forward, an innovation impacts y_t and y_{t-365} . For $h \geq 365$:

$$IRF_h^{YoY} = b_j = \phi_j - \phi_{j-365} \quad \text{for } j \geq 365. \quad (13)$$

Hence, the impulse response in levels for $h \geq 365$ can be retrieved from the YoY IRF

recursively according to

$$IRF_h = IRF_h^{YoY} + IRF_{h-365}. \quad (14)$$

In sum, the impulse response function in levels mapping to the YoY impulse response function is given by

$$IRF_h = \begin{cases} IRF_h^{YoY} & 0 \leq h < 365 \\ IRF_h^{YoY} + IRF_{h-365} & h \geq 365 \end{cases} \quad (15)$$

The year-on-year transformation also shapes the precision of the long-horizon projections. At horizon h , the dependent variable is $z_{t+h} = y_{t+h} - y_{t+h-365}$. Setting $j = 365 - h$, the lagged control $z_{t-j} = y_{t-j} - y_{t-j-365}$ has endpoint $y_{t-j} = y_{t+h-365}$ —the same base-year level that enters the dependent variable. With $p = 90$ daily lags, this overlap begins at $h = 275$ days and continues through $h = 364$. It gives the lagged controls predictive content for the long-horizon dependent variable and can raise the R^2 and lower the long-run variance entering the HAC standard error, and can account for some of the narrowing of the confidence intervals at long horizons. The overlap is with pre-shock variables, so it is not a bad-control channel and does not mechanically bias the monetary-shock coefficient, which is estimated from shock variation residualized with respect to the lagged controls under our maintained exogeneity condition. The smooth local projection leaves this indexing unchanged, though it is a separate source of regularization reflected in the reported standard errors. The narrower long-horizon bands are therefore best read as conditional, pointwise local-projection precision rather than as an artifact of the year-on-year-to-levels recovery.

B.3 Pre-Trend Test

The causal interpretation of local projection estimates requires the lead-lag exogeneity condition discussed in [Stock and Watson \(2018\)](#) and [Jordà and Taylor \(2025\)](#): the identified shock must be exogenous with respect to past outcomes. A necessary implication is that the shock should not be predictable from lagged outcomes and controls—that is, $\mathbb{E}[u_t \mid \mathcal{F}_{t-1}] = 0$, where $\mathcal{F}_{t-1}$ contains past outcomes and controls. If this condition fails, the estimated impulse responses may reflect anticipation or endogenous policy response rather than the causal effect of the shock. We test this necessary condition by extending the baseline local projection to negative horizons. Specifically, we estimate

$$y_{t+h} = \alpha_h + \beta_h \text{shock}_t + \theta_h \text{cases}_t + \delta_h \text{stringency}_t + \sum_{\ell=1}^p \varphi_{h,\ell} y_{t+h-\ell} + \varepsilon_{h,t} \quad (16)$$

for horizons $h = -60, -59, \dots, -1, 0, 1, \dots, 365$. This is the same specification as equation (1) in the main text, now evaluated at horizons that precede the shock. For $h < 0$, the dependent variable y_{t+h} is a realization of the outcome that has already occurred at time t . The lags of the dependent variable are timed relative to y_{t+h} (that is, $y_{t+h-1}, y_{t+h-2}, \dots, y_{t+h-p}$) rather than relative to the shock date, so that the dependent variable does not appear among the regressors. For $h \geq 0$, the specification coincides with the baseline. Standard errors are heteroskedasticity-robust (HC1), COVID observations (March 14 – October 30, 2020) are excluded, and coefficients are scaled by the standard deviation of the shock.


We estimate horizon-by-horizon OLS rather than smooth LP for this diagnostic. Smoothing across $h = 0$ would blend the pre-trend and post-shock regions; independent estimation at each horizon is the appropriate choice for testing whether individual pre-trend coefficients differ from zero.

The logic of the test is as follows. For $h < 0$, β_h measures whether a monetary policy surprise at date t predicts outcomes that have already been realized at $t + h < t$. Under valid identification, future shocks cannot cause past outcomes, so β_h should equal zero for all $h < 0$. Any systematic deviation would indicate either anticipation of the shock or endogenous response of monetary policy to pre-existing dynamics. Because the test uses the same specification as the baseline, the pre-trend coefficients are directly comparable to the post-shock impulse responses. The daily frequency of our data makes this test particularly informative: we observe 60 individual days before each shock, providing a level of granularity that is not available in settings where shocks and outcomes are measured at monthly or quarterly frequency.

Figure B2 reports the results for consumption, sales, employment, and investment. The left portion of each panel ($h < 0$, shaded) shows the pre-trend coefficients; the right portion ($h \geq 0$) shows the impulse response. For all four variables, the pre-trend coefficients are centered on zero and display no drift. The fraction of coefficients that reject zero at the 5% level ranges from 3% to 10% across variables, consistent with the nominal size of the test: under the null, we would expect 3 out of 60 coefficients (5%) to reject by chance, and the 95% binomial confidence interval spans roughly 1% to 13%. At $h = 0$, the response departs from zero and the break is not a continuation of any movement in the pre-shock window.

This diagnostic checks a necessary but not sufficient condition for lead-lag exogeneity. It establishes that the shock is not forecastable from observed past outcomes and controls, but does not rule out correlation with unobserved state variables or the condition $\mathbb{E}[u_t \varepsilon_{t+h}] = 0$ for $h > 0$, which is not directly testable. With this qualification, the results support the predeterminedness of the identified shock.

Figure B2: Pre-trend test: local projection coefficients at negative horizons



Notes: Each panel plots the horizon-by-horizon OLS coefficient $\hat{\beta}_h$ from equation (16) for $h = -60$ to $h = 365$ days. The shaded region ($h < 0$) corresponds to the pre-trend test: coefficients measure whether future monetary policy shocks predict past outcomes. The dashed line marks $h = 0$. Darker bands: 68% confidence intervals; lighter bands: 90% confidence intervals. Controls: log of new COVID cases, log of Oxford Stringency Index, and 90 lags of the dependent variable (timed before y_{t+h} for $h < 0$). COVID observations (March 14 – October 30, 2020) excluded. Coefficients scaled by the standard deviation of the monetary policy shock. Sample: August 2015 – October 2023.

C Appendix for Section 4

C.1 Empirical Evidence on High-Frequency Monetary Policy Shock Transmission: a Comparative Review of the Literature

This subsection contrasts our results regarding monetary transmission to financial and real variables with established literature, focusing on response magnitudes and dynamics across various settings and time periods.

Financial markets. Focusing on the monetary policy event window and using our baseline sample, the F-statistics from our first-stage regression—focusing on the 6-month and 1-year German yields in response to baseline monetary policy shocks—are 28.42 and 16.33, respectively. These outcomes are consistent with the literature, for instance, see the findings by [Altavilla et al. \(2019\)](#). When examining the data on a monthly basis, our results align both qualitatively and quantitatively with those of [Jarociński and Karadi \(2020\)](#) concerning the one-year German bond yield and the Euro-Area stock index, as depicted in Figure 8B of their study.

In comparison to the euro area, there is a much larger number of studies focused on the United States. Starting with high-frequency evidence, [Jarociński \(2024\)](#) estimates daily IRFs of US financial variables in response to monetary policy shocks. According to the findings by this author, the response of the stock market and inflation expectations are similar in magnitude to ours: around -0.5% in the first 25 days for the S&P500, close to our estimate of around -0.4% for IBEX35, and between -0.01% and -0.02% (depending on the shock considered) on impact for the 5-year breakeven inflation (TIPS), again in line with our estimates. For the 6-months US treasury yields, [Swanson \(2021\)](#) reports an IRF to an identified Federal funds rate shock that exhibits an increasing dynamic up to 120 days after the shock, similar to our estimated IRF for the euribor. Finally, [Lewis \(2023\)](#) finds that, empirically, the same day impact on financial variables may vary depending on the sample period and nature of the shock considered. For example, in the 1996-2019 sample, this author finds no significant impact of Fed funds rate shocks on the bond yields and the 10 years TIPS spreads. This is not the case (for this and other monetary policy shocks of different nature) in the sample 2009-2015.

Finally, at the monthly frequency, [Miranda-Agrippino and Ricco \(2021\)](#) report IRFs for the 1 year T-bond and the S&P500 similar in magnitude to our 12-months euribor and IBEX35 responses. After adjusting the magnitude of the shock for comparability, 6 months after the shock the response is about 0.015% for the 1 year T-bond and -0.18%

for the S&P500. The first estimate is close to our finding for the 12-months euribor, the second is of the same order of magnitude as the response of IBEX35, around -0.4% for IBEX35.

Real activity. To assess whether the real effects implied by our estimates are plausible, we compare our industrial production responses with those reported in the recent literature. Because our main analysis uses daily data, direct comparisons are limited. We therefore focus on industrial production (IP), which is available at monthly frequency and is commonly studied in related work. Figure G1 shows our responses estimated using standard, non-smoothed local projections. Two recent studies provide particularly relevant benchmarks: [Bauer and Swanson \(2023\)](#) for the United States and [Altavilla et al. \(2025\)](#) for the euro area. Both papers estimate the response of monthly IP to monetary policy shocks and, as in our analysis, use the high-frequency monetary policy shock directly as a regressor, rather than as an instrument in a proxy-SVAR or LP-IV framework.⁶⁹

To make the estimates comparable, we rescale the effects reported in [Bauer and Swanson \(2023\)](#) and [Altavilla et al. \(2025\)](#) to the size of the shock in our setting. The comparison yields two main takeaways. First, the magnitudes are broadly similar across studies. Our scaled trough response of industrial production is -0.72% , which lies between the corresponding estimates in [Bauer and Swanson \(2023\)](#) and [Altavilla et al. \(2025\)](#), equal to -0.89% and -0.53% , respectively. Second, the timing of the trough differs between the euro area and the United States. In line with [Altavilla et al. \(2025\)](#), we find that IP reaches its trough after about one month, whereas [Bauer and Swanson \(2023\)](#) report a trough around twelve months.⁷⁰ The close alignment with [Altavilla et al. \(2025\)](#) responses, whose euro-area setting is closer to ours, is therefore reassuring.

⁶⁹[Bauer and Swanson \(2023\)](#) estimate impulse responses using local projections, while [Altavilla et al. \(2025\)](#) use a Bayesian VAR.


⁷⁰[Bauer and Swanson \(2023\)](#) show that moving from local projections to a proxy-SVAR, or including Federal Reserve speeches, substantially affects the magnitudes of the responses. However, the timing and shape of the IP response are robust across specifications: industrial production reaches its trough after roughly 8 to 12 months in the United States regardless of the method used.

C.2 Cumulative Monetary Multipliers of Real Activity

Figure C1 presents the full cumulative monetary multiplier profiles underlying the summary statistics reported in the main text. Rather than focusing only on selected horizons, the figure traces the evolution of the cumulative response ratios over the full one-year window after the monetary policy shock.

The profiles show that the cumulative multipliers turn negative shortly after the shock across all outcomes. Sales and consumption exhibit a pronounced initial decline, followed by a gradual partial recovery over the medium horizon. Investment also falls on impact, though with wider uncertainty bands, reflecting the greater volatility of investment responses. Employment displays a smaller cumulative response in absolute value, but the decline is more persistent and remains clearly negative throughout the horizon.

Figure C1: Cumulative monetary multipliers of real activity



Notes: Each panel shows the cumulative response ratio of a real activity variable (in percentage) relative to the cumulative response of the 12-month Euribor (in basis points), following a one standard deviation contractionary monetary policy shock, providing an analogue to the cumulative fiscal multiplier of [Mountford and Uhlig \(2009\)](#) and [Ramey \(2016\)](#). Shaded areas denote 68% and 90% confidence intervals computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. The monetary policy shock standard deviation is 3.7 bp for consumption and employment and 4.1 bp for sales and investment.

C.3 The Response over a Two-Year Horizon

Our baseline projections run to a one-year horizon; here we extend them to two years, subject to a caveat about identification. The densest cluster of larger surprises in our sample falls in the European Central Bank’s 2022–2023 tightening cycle, and that cluster is mechanically truncated at long horizons: y_{t+h} is unavailable for shocks near the end of the sample, so the two-year coefficients are estimated from a thinner set of shocks than the short-horizon coefficients. Long-horizon projections therefore rely disproportionately on pre-2022 variation. The reported standard errors do not widen visibly, but they are conditional on this reduced long-horizon sample and do not capture the loss of support from the mechanically excluded 2022–2023 surprises. We accordingly treat the two-year magnitudes as exploratory.


Figure C2 extends the baseline daily projections. Sales and consumption reach their troughs within the first year and remain modestly negative thereafter, while employment and—more noisily—investment build more slowly, with a trough around eleven to twelve months before reverting in the second half of the second year. These estimates therefore confirm the pattern in our baseline, of early adjustment in final demand but slower responses in employment and investment. The longer horizon further suggests that the response of the latter two variables are broadly in line with the sixteen-month average lag that Friedman (1961) reports between peaks in the rate of change of the money stock and reference-cycle peaks.

In addition, we report two-year horizon projections on monthly data, to further inspect the response of slower-moving outcomes. Figure C3 reports two-year monthly responses for employment, the CPI and a broad measure of economic activity given by a monthly nowcast of GDP produced by The Spanish Independent Authority for Fiscal Responsibility (AIReF). The latter is available only for a short sample spanning December 2020–April 2025 rendering the results reported below tentative and subject to small sample noise.⁷¹

For this set of variables available at monthly frequency we find that employment declines gradually and remains negative through the second year; the GDP proxy traces a slow contraction that, in these estimates, deepens to about -0.5% near the thirteenth month before reverting; the CPI declines steadily, with a trough late in the horizon. The estimated troughs occur around one year or later. Whether the peak in the long-lag re-

⁷¹Specifically this measure is the AIReF’s *MIPred* real-time GDP nowcast, which we reconstruct into a continuous monthly series from archived copies of AIReF’s published vintages retrieved via the Internet Archive (Wayback Machine) and enter in log-levels in a plain local projection over December 2020–April 2025; the reconstruction is described in detail at the end of this appendix section (paragraph “Reconstructing the AIReF MIPred nowcast”).


Figure C2: Two-year response of real activity to a monetary policy shock



Notes: LP impulse response functions to a one standard deviation monetary policy shock, reported in levels, with the estimation horizon extended to two years. The monetary policy shock standard deviation is 4.1bp for sales and investment and 3.7bp for consumption and employment. Confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors; darker-shaded areas are 68% and lighter-shaded areas 90% confidence intervals. The sample ends October 30, 2023, and begins one year after the start of each series reported in Table A1.

sponse lies at one year, as some recent high-frequency LP studies suggest (Jarociński and Karadi, 2020; Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2023), at one to two years (Christiano, Eichenbaum and Evans, 1999), or at two to three years (Romer and Romer, 2004), is an open question our short sample—particularly for the GDP nowcast—cannot settle. We find peak effects in final demand much earlier than the classical literature, but our short sample cannot pin down where the long-lag peak in employment and prices falls; doing so would require longer samples and methods built for long-horizon estimation. The time-aggregation biases we document in Section 6 nonetheless suggests that late peaks found by earlier studies may partly reflect the quarterly frequency of the dat used in the estimation.

Figure C3: Two-year monthly responses of employment, GDP, and the CPI to a monetary policy shock




Notes: Monthly LP impulse responses, in levels, over a two-year horizon. Employment (daily employment aggregated to the monthly frequency) and the CPI (INE overall index) are year-on-year monthly local projections converted to levels by the recursion $L_h = \text{IRF}_h + L_{h-12}$, on the baseline sample (August 2015–October 2023); the GDP proxy is the AIReF MIPred real-time nowcast entered in log-levels over December 2020–April 2025. The three panels are estimated over different samples and are not comparable horizon by horizon. Confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors; darker-shaded areas are 68% and lighter-shaded areas 90% confidence intervals; given the short GDP sample, the GDP panel is exploratory.

Reconstructing the AIReF MIPred nowcast. MIPred is the dynamic-factor model that the Spanish Independent Authority for Fiscal Responsibility (AIReF) maintains to now-cast quarter-on-quarter real GDP growth from a panel of higher-frequency indicators. AIReF re-estimates it every few days, so each run produces a new vintage: the model’s estimate of every target quarter’s growth as of that date. Reading the successive vintages of a single quarter shows how its estimate is revised as new information arrives. The estimates AIReF currently publishes start only in January 2024, which leaves too short a real-time history for our purposes.

To go back further, we recover earlier releases from archived copies of AIReF’s website on the Internet Archive (Wayback Machine). Nine snapshots dated between September 2022 and June 2024 extend the vintage record back to December 2020; together with the current vintages they give 439 distinct vintages in all. Where two snapshots cover the same quarter at the same vintage date their estimates coincide, and that agreement is our internal check on the splicing. From these vintages we build a monthly real-time growth series with no gaps. To each calendar month we assign the most recent vintage dated on or before month-end and read its estimate of the current quarter. In the first month of a quarter, before any vintage has begun nowcasting it, we instead carry forward the most recent quarter already observed by INE.

We turn this growth series into a level in two steps. First, we compound the settled

Figure C4: The reconstructed MIPred real-time GDP series for Spain.



Panel (a) plots the monthly real-time quarter-on-quarter growth of real GDP read from the vintage panel; panel (b) plots the real-time GDP level index it implies, normalised to 100 in 2020Q3. The series is monthly and runs from December 2020 to May 2026.

quarter-on-quarter rates, each quarter's value as read from the latest available vintage, into a realised quarterly GDP index set to 100 in 2020Q3. Second, we fill in the within-quarter months, setting each month's level equal to the preceding quarter's realised index scaled by one plus that month's real-time growth rate, which we enter in log levels. Figure C4 plots the result: the real-time growth series in panel (a) and the level index it implies in panel (b). The series is monthly and runs from December 2020 to May 2026, and spans one re-specification of the model, the 2021 base being replaced by a 2024 base on 30 January 2024, with each vintage assigned to whichever specification was in production on its date. Two checks support the reconstruction. Across every closed quarter the settled growth rates match the published INE Quarterly National Accounts figures for quarter-on-quarter GDP growth exactly, and month-by-month spot checks reproduce the archived estimates. It is a model-based nowcast built from a panel of indicators rather than a direct measurement of activity.

D Appendix for Section 5

D.1 Upstream vs. Downstream Sectoral Classification

This section outlines how we bridge the sector classification used by the Spanish Tax Authority to compile their sales data with the 2015 Spanish INE Input-Output sector classification and how we calculate the upstreamness indicator and derive the sectoral upstream classification for the sales data.

Table D1 presents a mapping between the Spanish Tax Authority and the 2015 Spanish INE Input-Output sector classification. Out of the 64 sectors listed in the INE IO table, we can match 43 to the sales sector classification. Notably, health and education services remain unassigned, since there are no sales data recorded. Ultimately, we can link 20 sales sectors to the 43 sectors in the IO table, for which we can calculate an upstreamness indicator.

We adopt upstreamness metric proposed by Antràs et al. (2012), designed to accurately assess an industry position within the global production network. The upstreamness metric offers an insightful perspective on the relative distance of an industry from the final consumption phase, reflecting their involvement in early or intermediate stages in the supply chain. These authors show that this metric can be derived taking two distinct approaches, both anchored in input-output analysis.⁷² Both methods produce an upstreamness measure that is always at least one for every industry. Higher values indicate a greater degree of upstreamness for that particular industry.

Table D3 reports the upstreamness indicator for both the bridged sales sectors and those IO sectors that could not be linked to the sales data. The second column, labeled "In Sales Data?", reports whether that industry is part of the sales data with a simple "Yes" or "No" response. The last column, "Upstream", assigns a binary value (0 or 1) to indicate whether an industry is upstream (1)—upstreamness above the average of 2.20—or downstream (0)—below average. These findings align well with the results reported

⁷²One method calculates upstreamness by iterating through the production process to capture the weighted average position of an industry output in the overall value chain. Using input-output tables, the approach tracks the flow of goods and services across different industries, taking into account both direct and indirect uses of an industry output. Upstreamness is calculated by considering the number of times an industry's output is employed as an intermediate input by other industries prior to its ultimate consumption. The second approach involves using a system of equations. This technique characterizes upstreamness by utilizing a group of linear equations to express the interconnections between various industries. It operates on the principle that industries which supply a considerable amount of their production to other upstream industries should also be considered more upstream. By solving these equations, the method assigns a measure of upstreamness to each industry, indicating its role and significance in the initial stages of production.

Table D1: Summary of the Spanish Tax Authority Sectors and their corresponding row/columns in the IO Table

Sectors Spanish Tax Authority	Rows/Columns in IO Table
C10 + C11 + C12. Food Industry, Beverage and Tobacco Manufacturing	5
C13 + C14 + C15. Textile Industry, Apparel, and Footwear Manufacturing	6
C17 + C18. Paper Industry; Graphic Arts	8, 9
C20. Chemical Industry	11
C21. Pharmaceutical Products Manufacturing	12
C22 + C23. Rubber and Plastic Products Manufacturing; Other Non-Metallic Mineral Products Manufacturing	13, 14
C24 + C25. Metallurgy; Manufacturing of Metal Products, except Machinery and Equipment	15, 16
C26 + C27. Manufacturing of IT, Electronic and Optical Products; Manufacturing of Electrical Equipment	17, 18
C29. Manufacturing of Motor Vehicles, Trailers, and Semi-Trailers	20
C16 + C31. Wood and Cork Industry; Furniture Manufacturing	7, 22
C28 + C30 + C33. Machinery and Equipment Manufacturing n.e.c.; Other Transport Equipment Manufacturing; Repair and Installation of Machinery and Equipment	19, 21, 23
C19 + C32. Coke Ovens and Oil Refining; Other Manufacturing Industries	10
D. Electricity, Gas, Steam, and Air Conditioning Supply	24
F. Construction	27
G. Wholesale Trade and Trade Intermediaries, Repair of Motor Vehicles and Motorcycles and Retail Trade except Motor Vehicles and Motorcycles	28, 29, 30
H. Transport and Storage	31, 32, 33, 34
I. Hospitality	36
J. Information and Communications	35, 37, 38, 39, 40
M+N. Professional and Administrative Activities	45, 46, 47, 48, 49, 50, 51, 52, 53

Notes: The first column of this table displays the sectoral classification used by the daily sales weekly report of the Spanish Tax Authority. The second column shows which row and column in the INE IO table corresponds to the sales sectors. Each row/column in the IO table is a sector of activity that can be mapped directly to the Statistical classification of economic activities in the European Community (NACE rev. 2). This mapping is presented in Table D2.

by Antràs et al. (2012) for Spain. The upstreamness values presented in the table range from 1.01 for residential care services to 3.87 for mining and quarrying, demonstrating the variation in industries positions within the supply chain.

Table D2: Correspondence between Spanish Input-Output Table Sectors and NACE Rev. 2 Industry Classification

Rows/Columns in IO table	Industries	NACE rev. 2
1	Crop and animal production, hunting and related service activities	01
2	Forestry and logging	02
3	Fishing and aquaculture	03
4	Mining and quarrying	05–09
5	Manufacture of food products, beverages and tobacco products	10–12
6	Manufacture of textiles, wearing apparel and leather products	13–15
7	Manufacture of wood and of products of wood and cork, except furniture; manufacture of articles of straw and plaiting materials	16
8	Manufacture of paper and paper products	17
9	Printing and reproduction of recorded media	18
10	Manufacture of coke and refined petroleum products	19
11	Manufacture of chemicals and chemical products	20
12	Manufacture of basic pharmaceutical products and pharmaceutical preparations	21
13	Manufacture of rubber and plastic products	22
14	Manufacture of other non-metallic mineral products	23
15	Manufacture of basic metals	24
16	Manufacture of fabricated metal products, except machinery and equipment	25
17	Manufacture of computer, electronic and optical products	26
18	Manufacture of electrical equipment	27
19	Manufacture of machinery and equipment n.e.c.	28
20	Manufacture of motor vehicles, trailers and semi-trailers	29
21	Manufacture of other transport equipment	30
22	Manufacture of furniture; other manufacturing	31–32
23	Repair and installation of machinery and equipment	33
24	Electricity, gas, steam and air conditioning supply	35
25	Water collection, treatment and supply	36
26	Sewerage; waste collection, treatment and disposal activities; materials recovery; remediation activities and other waste management services	37–39
27	Construction	41–43
28	Wholesale and retail trade and repair of motor vehicles and motorcycles	45
29	Wholesale trade, except of motor vehicles and motorcycles	46
30	Retail trade, except of motor vehicles and motorcycles	47
31	Land transport and transport via pipelines	49
32	Water transport	50
33	Air transport	51
34	Warehousing and support activities for transportation	52
35	Postal and courier activities	53
36	Accommodation, food and beverage service activities	55–56
37	Motion picture, video and television programme production, sound recording and music publishing activities; programming and broadcasting activities	59–60
38	Telecommunications	61
39	Computer programming, consultancy and related activities; information service activities	62–63
40	Financial service activities, except insurance and pension funding	64
41	Insurance, reinsurance and pension funding, except compulsory social security	65
42	Activities auxiliary to financial services and insurance activities	66
43	Real estate activities	68
44	Imputed rent of owner-occupied dwellings	68a
45	Legal and accounting activities; activities of head offices; management consultancy activities	69–70
46	Architectural and engineering activities; technical testing and analysis	71
47	Scientific research and development	72
48	Advertising and market research	73
49	Other professional, scientific and technical activities; veterinary activities	74–75
50	Rental and leasing activities	77
51	Employment activities	78
52	Travel agency, tour operator reservation service and related activities	79
53	Security and investigation activities; services to buildings and landscape activities; office administrative, office support and other business support activities	80–82
54	Public administration and defence; compulsory social security	84
55	Education	85
56	Human health activities	86
57	Social work activities	87–88
58	Creative, arts and entertainment activities; libraries, archives, museums and other cultural activities; gambling and betting activities	90–92
59	Sports activities and amusement and recreation activities	93
60	Activities of membership organisations	94
61	Repair of computers and personal and household goods	95
62	Other personal service activities	96
63	Activities of households as employers; undifferentiated goods- and services-producing activities of households for own use	97–98
64	Activities of extraterritorial organisations and bodies	99


Table D3: Upstreamness Indicator and Upstream vs. Downstream Classification of Bridged IO Sectors

Industry	In Sales Data?	Upstreamness	Upstream
Residential care services; social work services without accommodation	No	1.01	0
Public administration and defense services; compulsory social security	No	1.06	0
Human health services	No	1.20	0
Education services	No	1.20	0
Other personal services	No	1.24	0
I. Hospitality	Yes	1.28	0
Creative, arts and entertainment services; libraries, archives, museums	No	1.42	0
Fish and other fishing products; aquaculture products	No	1.45	0
C21. Pharmaceutical Products Manufacturing	Yes	1.55	0
F. Construction	Yes	1.56	0
Sporting services and amusement and recreation activities	No	1.57	0
Real estate services	No	1.58	0
Services furnished by membership organizations	No	1.76	0
G. Wholesale Trade and Trade Intermediaries, Sale and Repair of Motor Vehicles and Motorcycles and Retail Trade, except Motor Vehicles and Motorcycles	Yes	1.86	0
Insurance, reinsurance, and pension funding services, except compulsory social security	No	1.89	0
Repair services of computers and personal and household goods	No	1.89	0
C13 + C14 + C15. Textile Industry, Apparel, and Footwear Manufacturing	Yes	2.02	0
C26 + C27. Manufacturing of IT, Electronic and Optical Products; Manufacturing of Electrical Equipment	Yes	2.09	0
C29. Manufacturing of Motor Vehicles, Trailers, and Semi-Trailers	Yes	2.16	0
Manufacturing of machinery and equipment n.e.c.; other transport equipment manufacturing; repair and installation of machinery and equipment	No	2.17	0
Natural water; water treatment and supply services	No	2.18	0
C16 + C31. Wood and Cork Industry; Furniture Manufacturing	Yes	2.19	0
J. Information and Communications	Yes	2.20	0
C10 + C11 + C12. Food Industry, Beverage and Tobacco Manufacturing	Yes	2.27	1
Services auxiliary to financial services and insurance	No	2.39	1
Financial services, except insurance and pension funding	No	2.58	1
C19 + C32. Coke Ovens and Oil Refining; Other Manufacturing Industries	Yes	2.65	1
Products of agriculture, hunting, and related services	No	2.76	1
M+N. Professional and Administrative Activities	Yes	2.78	1
H. Transport and Storage	Yes	3.09	1
Sewerage services; sewage sludge; waste collection and disposal services	No	3.11	1
D. Electricity, Gas, Steam, and Air Conditioning Supply	Yes	3.28	1
C22 + C23. Rubber and Plastic Products Manufacturing; Other Non-Metallic Mineral Products Manufacturing	Yes	3.31	1
Products of forestry, logging, and related services	No	3.33	1
C24 + C25. Metallurgy; Manufacturing of Metal Products, except Machinery and Equipment	Yes	3.53	1
C17 + C18. Paper Industry; Graphic Arts	Yes	3.55	1
Chemical industry	No	3.60	1
Mining and quarrying	No	3.87	1

D.2 Raw (Unsmoothed) Seasonally-Adjusted Local-Projection Overlays


The baseline daily impulse responses reported in the main text—for financial markets (Figure 1), consumption categories (Figure 3), sectoral sales (Figure 4), and the upstream vs. downstream panel (Figure 5)—are estimated with the smooth local projection of [Barnichon and Brownlees \(2019\)](#), as described in Section 3.3. Following Referee 1’s suggestion, and for transparency, we reproduce each of these figures below with the corresponding *standard, unsmoothed* local projection overlaid in grey (“Regular LP”); the red line and blue bands are the smooth-LP baseline (“Smooth LP”). For the real-activity figures the overlaid local projection is estimated on the seasonally adjusted series; for the financial-market series—which are not seasonally adjusted—it is the standard local projection on the same raw series, estimated with the same controls as the smooth baseline, so that the two estimators differ only in the smoothing step. In the main text we plot the raw overlay only for the four headline activity variables (Figure 2) and report the smooth baseline alone for these remaining daily series, to keep the figures legible. The overlays collected here make transparent that the short-lag signal is already visible in the noisier raw estimates—before any smoothing—and that the smooth baseline traces the center of the raw responses rather than creating the signal.

Figure D1: Financial markets: smooth-LP baseline with standard (unsmoothed) LP overlay




Notes: As Figure 1, with the corresponding standard, unsmoothed local projection on the same (raw, not seasonally adjusted) series overlaid in grey ("Regular LP"); the red line and blue bands are the smooth-LP baseline ("Smooth LP"). Both estimators use the same controls and differ only in the smoothing step. Confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors; darker-shaded areas are 68% and lighter-shaded areas 90% confidence intervals. The sample starts in August 1st, 2015 and ends in October 30th, 2023. The monetary policy shock standard deviation is 3.7bp.

Figure D2: Consumption by category: smooth-LP baseline with raw seasonally-adjusted overlay




Notes: As Figure 3, with the standard (unsmoothed) local projection on the seasonally adjusted data overlaid in grey (“Regular LP”); the red line and blue bands are the smooth-LP baseline (“Smooth LP”). Confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors; darker-shaded areas are 68% and lighter-shaded areas 90% confidence intervals. The sample starts in August 1st, 2016 and ends in October 30th, 2023. The monetary policy shock standard deviation is 3.7bp.

Figure D3: Sales by sector: smooth-LP baseline with raw seasonally-adjusted overlay



Notes: As Figure 4, with the standard (unsmoothed) local projection on the seasonally adjusted data overlaid in grey (“Regular LP”); the red line and blue bands are the smooth-LP baseline (“Smooth LP”). Confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors; darker-shaded areas are 68% and lighter-shaded areas 90% confidence intervals. The sample starts in July 1st, 2018 and ends in October 30th, 2023. The monetary policy shock standard deviation is 4.1bp.

Figure D4: Upstream vs. downstream sectoral sales: smooth-LP baseline with raw seasonally-adjusted overlay



Notes: As Figure 5, with the standard (unsmoothed) local projection on the seasonally adjusted panel overlaid in grey (“Regular LP”); the colored lines and bands are the smooth-LP baseline (“Smooth LP”), with upstream the solid red and downstream the dashed blue line. The left panel reports the upstream and downstream responses; the right panel their difference. The confidence intervals are computed from heteroskedasticity-robust standard errors, clustered by time. Shaded areas are the 90% confidence intervals (left panel) and the 68% and 90% confidence intervals (right panel) of the smooth-LP baseline. The sample starts in July 1st, 2018 and ends in October 30th, 2023. The monetary policy shock standard deviation is 4.1bp.

E Appendix for Section 6


E.1 Comparison of Lower-Frequency LP IRFs with Baseline Daily IRFs at Daily Frequency

In the main text, we contrasted IRFs derived from low-frequency data with those from time-aggregated daily data LP IRFs, effectively situating all IRFs within a low-frequency context. Alternatively, IRFs estimated from LPs using low-frequency data can be combined into a mixed-frequency dataset, presenting all estimated IRFs at the daily frequency. Although comparing point estimates over the same period from different frequencies becomes visually challenging, it enhances the ability to appreciate the superior resolution offered by high-frequency estimates. In this section, we provide this alternative comparison for weekly, monthly, and quarterly frequencies in Figures E1, E2 and E3, respectively. Examining these figures, we derive conclusions similar to those in the main text. The IRFs of daily data aggregated into weekly and monthly frequencies closely follow the daily IRFs, whereas the quarterly data differ. In the first quarter, responses of sales, employment and consumption become insignificant while they were detectable at the weekly and monthly frequencies. Moreover, the comparison underscores the importance of preserving the high resolution of the daily data. For investment, lower frequencies produce insignificant short lags in all low-frequency estimates, whereas our high-resolution daily IRFs detect significant responses at a confidence level 90% around 45 days after monetary policy shock. Furthermore, high resolution reveals that significant responses in sales, consumption, and employment typically materialize at the month's end, implying that the anticipated contemporaneous impact in monthly studies depends on whether the shock occurs at the start or end of the month.

E.2 Simulation exercise


This appendix details the simulation designs described in the main text and reports further results. The logic of the exercise is simple. We construct a daily economy in which the true impulse response to a monetary policy shock is, by construction, the one we estimate in the data. We then observe this economy as a typical researcher would—at monthly or quarterly frequency—and ask whether standard local projections recover the truth. Because we control the data-generating process, any gap between the estimates and the truth is estimation bias, and we can trace it to its source. We proceed in four steps: we first describe the data-generating process and its calibration; we then document

Figure E1: Time Aggregation: Weekly responses




Notes: LP impulse response functions to a one standard deviation monetary policy shock. The shock size is identical across variables and frequencies: all responses are scaled by the standard deviation of the daily shock series, as in our baseline. The responses are reported in levels. The weekly estimates are standard (unsmoothed) local projections with the COVID treatment of Section 3.3, on weekly aggregates of the seasonally adjusted daily series of Section 3.2. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in October 30th, 2023. The sample starts one year after the sample start of each series reported in Table A1. Clockwise, we display LP for total sales, employment, consumption and investment. Solid step lines are the low-frequency aggregated data (weekly) LP point estimates, plotted at daily resolution; dashed lines are the baseline daily LP point estimates.

Figure E2: Time Aggregation: Monthly responses



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The shock size is identical across variables and frequencies: all responses are scaled by the standard deviation of the daily shock series, as in our baseline. The responses are reported in levels. The monthly estimates are standard (unsmoothed) local projections with the COVID treatment of Section 3.3, on monthly aggregates of the seasonally adjusted daily series of Section 3.2. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in October 30th, 2023. The sample starts one year after the sample start of each series reported in Table A1. Clockwise, we display LP for total sales, employment, consumption and investment. Solid step lines are the low-frequency aggregated data (monthly) LP point estimates, plotted at daily resolution; dashed lines are the baseline daily LP point estimates.

Figure E3: Time Aggregation: Quarterly responses



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The shock size is identical across variables and frequencies: all responses are scaled by the standard deviation of the daily shock series, as in our baseline. The responses are reported in levels. The quarterly estimates are standard (unsmoothed) local projections with the COVID treatment of Section 3.3, on quarterly aggregates of the seasonally adjusted daily series of Section 3.2. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in October 30th, 2023. The sample starts one year after the sample start of each series reported in Table A1. Clockwise, we display LP for total sales, employment, consumption and investment. Solid step lines are the low-frequency aggregated data (quarterly) LP point estimates, plotted at daily resolution; dashed lines are the baseline daily LP point estimates.

the bias that time aggregation produces at monthly and, especially, quarterly frequency; we extend the simulated sample to assess whether the bias merely reflects the relatively short quarterly sample; and, finally, we isolate the source of the bias with designs that vary the timing of shocks and the shape of the impulse response function in turn.

Simulation of Daily Data Let g_t be the year-on-year log growth rate (in percent) of an economic variable of interest on day t , matching the outcome variable in our local projections. Each day t is associated with a monetary policy shock s_t . As a parsimonious model of the persistence in daily growth not explained by monetary policy shocks, we use the following AR(1) model⁷³ as a data-generating process relating growth to shocks:

$$g_t = \alpha + \rho g_{t-1} + \sum_{h=0}^{365} \theta_h s_{t-h} + \gamma^T \mathbf{X}_t + \epsilon_t \quad (17)$$

where $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$ is i.i.d. and independent of the monetary policy shocks at all leads and lags, and $\mathbf{X}_t = (\text{Cases}_t, \text{Stringency}_t)$ collects the COVID controls described in the main text.

Conditional on $\mathbf{X}_t$ and the independence of past shocks and s_t , the impact of s_t on g_t is θ_0 . The impact of s_t on g_{t+1} instead has two components. The first is a direct effect θ_1 . The second is an indirect effect $\rho\theta_0$, which arises because s_t affects g_{t+1} via its effect on g_t . More generally, given $\mathbf{X}_t$,

$$\frac{dg_{t+h}}{ds_t} = \sum_{i=0}^h \theta_i \rho^{h-i} = \theta_h + \rho \sum_{i=0}^{h-1} \theta_i \rho^{h-1-i} = \theta_h + \rho \frac{dg_{t+h-1}}{ds_t} \quad (18)$$

In words, a shock affects growth both directly at each horizon and indirectly through the persistence of growth itself. Now suppose one fits a local projection on data generated from (17) of the form

$$g_{t+h} = \alpha_h + \beta_h s_t + \delta^T \mathbf{X}_t + \epsilon_t.$$

The estimated $\hat{\beta}_h$ captures the total effect of the shock on g_{t+h} , i.e. $\frac{dg_{t+h}}{ds_t}$. If one wishes (17) to satisfy $\frac{dg_{t+h}}{ds_t} = \hat{\beta}_h$, then θ_h must equal $\hat{\beta}_h - \rho\hat{\beta}_{h-1}$ by (18), with the convention $\hat{\beta}_{-1} = 0$ so that $\theta_0 = \hat{\beta}_0$. In particular, let $\hat{\beta}'_h$ be the smooth local projection coefficient for consumption growth from Figure 2. We convert this from standard deviation units into basis point units to obtain $\hat{\beta}_h = \hat{\beta}'_h / \sigma_{MP}^\epsilon$, which we use to set θ_h ; here σ_{MP}^ϵ is the standard

⁷³We also considered richer AR(p) specifications but use the AR(1) as a parsimonious benchmark for the component of daily growth that is independent of the simulated monetary policy shocks by construction. For any autoregressive specification, the distributed-lag coefficients can be recalibrated so that the daily-growth IRF equals the empirical IRF at every modeled horizon.

deviation of the daily monetary policy shock series, computed over all days from August 1, 2015 through June 15, 2023—including days without a policy announcement, on which the shock is zero. In other words, (17) is specified so that the total effect of a shock on growth is consistent with our baseline consumption growth IRF.

The remaining parameters of (17) govern how growth behaves absent monetary policy. To estimate them, note that we can rewrite the DGP as

$$g_t = \alpha + \rho g_{t-1} + \sum_{h=0}^{365} (\hat{\beta}_h - \rho \hat{\beta}_{h-1}) s_{t-h} + \gamma^T \mathbf{X}_t + \epsilon_t \implies$$

$$g_t - \sum_{h=0}^{365} \hat{\beta}_h s_{t-h} = \alpha + \rho \left(g_{t-1} - \sum_{h=0}^{365} \hat{\beta}_{h-1} s_{t-h} \right) + \gamma^T \mathbf{X}_t + \epsilon_t \quad (19)$$

which, up to a truncation term in $\hat{\beta}_{365}$, defines an AR(1) process over the residuals of g_t after removing the effect of monetary policy. We generate consumption growth residuals in this way from October 29, 2016 through June 15, 2023 (from the start of the local projection estimation sample through the last observed shock) using observed monetary policy shocks. We then fit (19) by OLS. The estimation sample includes the COVID-19 period: we control for COVID dynamics through $\mathbf{X}_t$ rather than dropping observations, in contrast to the baseline local projections. Table E1 reports the fitted values.

Table E1: Fitted Values for Consumption Growth Process

α	ρ	σ	γ_{Cases}	$\gamma_{\text{Stringency}}$
3.098	0.299	8.559	0.180	-0.862

The simulated economy also requires shocks that resemble those in the data. To obtain them, we fit a student-t distribution to the 63 shocks observed from August 1, 2015 through June 15, 2023—the period we use to obtain the scaling factor σ_{MP}^ϵ . We fix the degrees of freedom at three to account for heavy tails, and obtain the location (0.89) and scale (2.42) by matching first and second moments. Let S denote the resulting distribution.

With the growth process and the shock distribution in hand, we simulate 1,000 daily consumption growth series from August 1, 2015 through June 15, 2023 in the following way:⁷⁴

1. For each day t with a non-zero shock in the real data, draw $s_t \sim S$, the calibrated student-t distribution of monetary policy shocks described above, independently

⁷⁴The estimation sample starts on October 29, 2016 but we begin the simulation from August 1, 2015 to allow for a burn-in period to reduce the impact of the initial condition.

across days and simulations. For each day with a zero shock in the real data, and for all days before August 1, 2015 entering the lag sum, set $s_t = 0$.

2. As the initial condition, draw the first day's growth uniformly at random from an interval centered on the unconditional mean of the shock-free autoregressive component: $g_1 \sim U[\hat{\mu} - 0.5, \hat{\mu} + 0.5]$ with $\hat{\mu} = \frac{\hat{\alpha}}{1-\hat{\rho}}$.
3. For each of the subsequent days, draw $g_t \sim \mathcal{N}(\hat{\alpha} + \hat{\rho}g_{t-1} + \sum_{h=0}^{365} \theta_h s_{t-h}, \hat{\sigma}^2)$. We set the COVID controls to zero, thereby removing their modeled contribution from the simulated paths.

Aggregating Daily Data

We now take the position of a researcher who observes this economy only at lower frequency. We aggregate the daily data to a given frequency and estimate local projections via OLS, measuring the monetary policy shock as the sum of daily shocks within the period. In line with the main text, quarterly estimation includes one lag of the dependent variable and two lagged shocks as covariates; monthly estimation includes three lags of the dependent variable and six lagged shocks.

Aggregation requires levels: quarterly consumption is the sum of daily consumption within the quarter, not the sum of daily growth rates. Because we simulate growth rates, we must therefore first recover levels. Since g_t is a daily year-on-year log growth rate, it satisfies

$$g_t = 100 \times (\log c_t - \log c_{t-365})$$

where $t - 365$ denotes the calendar date exactly 365 days earlier, for t from 2016-07-31 through 2023-06-15. To recover daily consumption *levels*, we use observed levels c_t for the first 365 days of the sample (2015-08-01 through 2016-07-30); for dates from 2016-07-31 onward, we invert the log growth definition. Rearranging the above equation yields

$$c_t = c_{t-365} \times \exp(g_t/100).$$

We apply this formula sequentially: for 2016-07-31, we use the observed level from 2015-08-01 and the simulated growth rate for 2016-07-31; for 2016-08-01, we use the observed level from 2015-08-02 and the simulated growth rate for 2016-08-01; and so on. From 2017-07-31 onward, the lagged level c_{t-365} is itself a previously computed simulated level rather than an observed one. We then aggregate to lower frequency (monthly and quarterly) by summing daily levels within each period and computing year-on-year growth rates.

Time Aggregation Bias

We now define precisely time aggregation bias. Consider quarterly aggregation (we use analogous objects for monthly-aggregated data). The target value is the average value at quarterly frequency of the daily IRF used to simulate consumption growth data. At horizon Q0 this is $\frac{1}{91} \sum_{h=0}^{90} \hat{\beta}'_h$, at horizon Q1 $\frac{1}{91} \sum_{h=91}^{181} \hat{\beta}'_h$, and so on. Notice that we express the target in standard deviation units rather than basis units to align with the scale from Figure 2. The estimated IRF is the average value of the local projection coefficients across the 1,000 simulations. We scale the coefficients, and their standard errors, by σ_{MP}^ϵ so they are comparable in scale to the target. Time aggregation bias is the difference between the estimate and target. In addition, we also report coverage—the relevant quantity for inference—of the 68% and 90% confidence intervals. This is the fraction of simulations in which confidence intervals of the given size contain the target. Well-calibrated intervals have coverage close to their nominal level (68% or 90%).

Figure E4 compares the estimated IRF against the target at monthly and quarterly frequencies. Some bias appears at monthly frequency, but the estimated pattern qualitatively matches that of the Consumption IRF aggregated to monthly frequency. At quarterly frequency, in contrast, the bias is large at the first (Q0) and second (Q1) horizons, and the estimated dynamics no longer match those of the Consumption IRF: the response at Q1 is estimated as larger than at Q0, whereas the former is 4.6 times smaller in the Consumption IRF. Quarterly aggregation does not merely attenuate the short-lag response; it moves the estimated trough from Q0 to Q1, mirroring the delayed profile in Section 6, where quarterly estimation fails to detect a statistically significant same-quarter consumption response.


Table E2 reports bias and coverage at all horizons. Monthly estimation produces biases at short horizons that largely dissipate by M4. Quarterly estimation fares much worse: coverage for every quarter is worse than for any individual month in the same quarter, and the biases at Q0, Q1, and Q2 are larger than at any corresponding monthly horizon.

Time Aggregation Bias with Extended Sample


A natural concern is that these results reflect the short quarterly sample rather than aggregation itself. To address it, we extend the simulation of daily data through December 31, 2043, bringing the number of quarters from 25 (2017Q1 through 2023Q1 in the baseline) to 108 (2017Q1 through 2043Q4). The first step is to determine the dates on which to draw shocks. Through 2023Q2, we use the dates in our actual data. From then onwards, we simulate the days on which a non-zero shock occurs. Over 2019-2022, the

Figure E4: Time-Aggregated Consumption Impulse Response Functions

(a) Monthly



(b) Quarterly



Notes: The black curves plot the average values of the daily Consumption IRF displayed in Figure 2 at monthly (top panel) and at quarterly (bottom panel) frequency. The blue dashed curve displays the average IRF estimated from aggregated daily data across 1,000 simulations. The dark and light colored ranges indicate the ranges in which 68% and 90% of the point estimates fall across simulations.

average days of the year with a non-zero shock are (after rounding) 26, 69, 109, 158, 202, 253, 300, and 348. To generate non-zero shocks in a year, we sample a day in a seven-day window around each reference day. For example, the first shock in a year occurs on a randomly selected day from 23 through 29. On the selected day, we draw the value of the shock from the same calibrated student-t distribution as in the baseline. This randomization around reference dates preserves the property that each quarter has exactly two shocks. Given simulated shocks through December 31, 2043, we generate consumption growth with the same AR(1) process as above, recover consumption levels with the same chaining approach, and conduct the analysis at both monthly and quarterly frequency.

Table E3 shows the results. The bias is similar to that in the baseline, but coverage deteriorates sharply at the short, biased horizons: the longer sample shrinks standard errors, while point estimates remain centered away from the true values. More data thus yields precision around biased estimates, not convergence to the true values.

Time Aggregation Bias and Shock Timing

Another possibility is that the timing of shocks within periods drives inference distortions. Suppose that in each low-frequency period of X days there is one shock that occurs on the first day. Then the estimated low-frequency IRF at the shortest horizon will capture the behavior of the daily IRF over its first X days, and so on through the remaining horizons. Now suppose instead that there is one shock that occurs near the end of each period. Then the low-frequency IRF at the shortest horizon will only capture the part of the daily IRF that falls within the period rather than the full X -day response. When the initial response of the daily IRF differs from the full X -day response, there will be a bias.

In our actual data, there are exactly two shocks per quarter, one occurring early and one occurring late. In order to assess whether this timing contributes to inference distortion, we conduct three additional simulation exercises within the long series setting, each with a single shock per quarter. In the first, the shock is placed on the first calendar day of each quarter, so that it is perfectly aligned with the aggregation window. In the second, the shock occurs near the location of the actual first shock in each quarter, on average about 68 days before quarter-end. In the third, it occurs near the location of the second shock, on average about 20 days before quarter-end; recall that the consumption IRF troughs at 44 days. The remaining details remain the same as in the baseline.

Figure E5 shows the results, together with the two-shock setting for comparison. The first-day-only setting—in which the shock is perfectly aligned with the aggregation window—recovers the IRF assumed in the DGP essentially perfectly. This is the sharpest

Table E2: Time-Aggregated Consumption IRF Estimation: Bias and Coverage

(a) Monthly						(b) Quarterly					
Horizon	Target	Monthly Estimation				Horizon	Target	Quarterly Estimation			
		Mean	Bias	Coverage				Mean	Bias	Coverage	
				68%	90%					68%	90%
M0	-0.200	-0.066	0.134	19.1	40.3	Q0	-0.281	-0.114	0.167	10.2	23.8
M1	-0.376	-0.321	0.056	53.6	77.0	Q1	-0.061	-0.185	-0.124	19.7	36.0
M2	-0.269	-0.318	-0.049	57.0	80.1	Q2	-0.142	-0.063	0.079	38.4	62.4
M3	-0.108	-0.151	-0.042	59.1	80.5	Q3	-0.232	-0.198	0.034	58.0	79.1
M4	-0.038	-0.035	0.003	63.0	86.6						
M5	-0.041	-0.020	0.021	65.1	87.0						
M6	-0.076	-0.049	0.026	62.4	84.0						
M7	-0.129	-0.094	0.035	60.7	83.0						
M8	-0.206	-0.162	0.044	58.1	81.7						
M9	-0.262	-0.228	0.034	63.3	84.1						
M10	-0.257	-0.241	0.016	64.3	84.8						
M11	-0.189	-0.187	0.002	64.7	87.1						

Notes: The ‘target’ column displays the average values of the daily consumption IRF displayed in Figure 2 at monthly (left panel) and at quarterly (right panel) frequency. ‘Mean’ is the average value of the point estimates from local projections fit on aggregated daily data across 1,000 simulations. ‘Coverage’ is the fraction of simulations in which the target falls in confidence intervals of a given size.

Table E3: Time-Aggregated Consumption IRF Estimation: Bias and Coverage (Extended Sample)

(a) Monthly

Horizon	Target	Monthly Estimation			
		Mean	Bias	Coverage	
				68%	90%
M0	-0.200	-0.068	0.132	0.4	1.9
M1	-0.376	-0.331	0.045	40.2	66.2
M2	-0.269	-0.339	-0.069	20.8	43.4
M3	-0.108	-0.171	-0.063	30.2	51.3
M4	-0.038	-0.056	-0.018	61.3	84.4
M5	-0.041	-0.031	0.010	64.4	86.6
M6	-0.076	-0.056	0.019	62.0	84.5
M7	-0.129	-0.106	0.023	55.7	79.9
M8	-0.206	-0.177	0.029	54.0	78.5
M9	-0.262	-0.243	0.018	61.9	83.9
M10	-0.257	-0.269	-0.012	62.4	86.9
M11	-0.189	-0.206	-0.017	61.3	88.2


(b) Quarterly

Horizon	Target	Quarterly Estimation			
		Mean	Bias	Coverage	
				68%	90%
Q0	-0.281	-0.128	0.153	0.4	1.7
Q1	-0.061	-0.195	-0.134	0.1	0.7
Q2	-0.142	-0.074	0.068	12.3	28.8
Q3	-0.232	-0.216	0.016	56.8	80.7

Notes: We document the performance of local projection estimation on an extended sample of simulated daily data through December 31, 2043. The ‘Target’ column displays the average values of the daily Consumption IRF displayed in Figure 2 at monthly (left panel) and at quarterly (right panel) frequency. ‘Mean’ is the average value of the point estimates from local projections fit on aggregated daily data across 1,000 simulations. ‘Coverage’ is the fraction of simulations in which the target falls in confidence intervals of a given size.

evidence on the source of the bias: once the shock is aligned with the window, the bias disappears. The first-shock setting still exhibits some bias, but tracks the true IRF much better than the two-shock setting. The second-shock setting performs worst, particularly at the first two horizons. When the shock arrives late in the quarter, the quarterly regression misses the response that has not yet materialized within the quarter and attributes the delayed peak to the next one. Table E4 documents the corresponding bias and coverage.

Figure E5: Quarterly IRFs with Alternative Shock Structures (Extended Sample)



Notes: We estimate quarterly IRFs on four different simulated time series data from 2017Q1 through 2043Q4. The first (“Two shocks”) is the extended-sample baseline, with two shocks per quarter as in the data. The second (“First-day only”) draws one shock per quarter on the first day of the quarter. The third (“First shock only”) draws one shock per quarter at a point similar to the location of first-occurring shocks in our data. The fourth (“Second shock only”) draws one shock per quarter at a point similar to the location of second-occurring shocks in our data. The graph shows average point estimates across 1,000 simulations compared to the quarterly average of the IRF in the DGP (“Consumption IRF”).

Time Aggregation Bias and IRF Shape

As a final exercise, we replace the consumption growth IRF with the employment growth IRF in the data-generating process. Specifically, we use (17) to simulate g_t by setting $\theta_h = (\hat{\beta}'_h - \hat{\rho}\hat{\beta}'_{h-1})/\sigma_{MP}^\epsilon$ where $\hat{\beta}'_h$ is the Employment IRF at horizon h from Figure 2 and $\hat{\rho} = 0.299$. We continue to use the α , ρ , and σ fit on consumption growth to isolate the impact of changing the IRF shape alone.

Unlike consumption growth, which reacts to shocks by declining quickly before re-

Table E4: Quarterly IRF Estimation with Alternative Shock Structures: Bias and Coverage

Horizon	Target	First-day only			First shock only			Second shock only		
		Bias	Coverage		Bias	Coverage		Bias	Coverage	
			68%	90%		68%	90%		68%	90%
Q0	-0.281	-0.006	61.9	83.7	0.056	23.5	43.4	0.246	0.0	0.0
Q1	-0.061	-0.001	65.3	87.6	-0.051	26.2	47.3	-0.221	0.0	0.0
Q2	-0.142	-0.002	64.1	86.2	0.040	37.0	63.2	0.095	5.8	16.8
Q3	-0.232	-0.004	63.8	84.9	-0.015	59.3	83.2	0.041	39.3	63.6

Notes: We estimate quarterly IRFs on three different simulated time series data from 2017Q1 through 2043Q4. The first (“First-day only”) draws one shock per quarter on the first day of the quarter. The second (“First shock only”) draws one shock per quarter at a point similar to the location of first-occurring shocks in our data. The third (“Second shock only”) draws one shock per quarter at a point similar to the location of second-occurring shocks in our data. The reports bias and coverage from 1,000 simulations compared to the quarterly average of the Consumption IRF in the DGP (“Target”).

covering, employment growth reacts with a steadily accumulating negative effect. This allows us to explore whether a distinct IRF shape eliminates aggregation bias. Table E5 (Table ?? in the main paper) reports results. Across all quarterly horizons, we observe an attenuation bias of varying size. On impact the bias is notably large: in percentage terms, the Q0 bias is actually more severe than for consumption (+0.059 against a target of -0.082). Aligning a single shock with the first day of each quarter again makes estimation almost unbiased (third panel), confirming that shock timing is responsible for aggregation bias.

Table E5: Time-Aggregated Employment IRF Estimation: Bias and Coverage

(a) Monthly						(b) Quarterly					
Horizon	Target	Monthly Estimation				Horizon	Target	Quarterly Estimation			
		Mean	Bias	Coverage				Mean	Bias	Coverage	
				68%	90%					68%	90%
M0	-0.042	-0.008	0.034	58.3	78.9	Q0	-0.082	-0.023	0.059	38.3	61.8
M1	-0.094	-0.068	0.026	61.6	84.0	Q1	-0.105	-0.102	0.004	64.9	84.8
M2	-0.109	-0.097	0.012	63.5	85.3	Q2	-0.134	-0.113	0.022	58.2	79.9
M3	-0.109	-0.101	0.008	66.0	86.9	Q3	-0.190	-0.150	0.041	50.3	72.2
M4	-0.103	-0.096	0.007	61.2	83.4						
M5	-0.104	-0.097	0.007	63.3	85.9	(c) Quarterly, first-day-only shocks					
M6	-0.120	-0.111	0.009	62.8	85.5						
M7	-0.133	-0.124	0.010	63.5	85.5						
M8	-0.146	-0.127	0.019	61.8	83.8						
M9	-0.167	-0.148	0.019	61.6	83.7						
M10	-0.189	-0.167	0.021	60.9	82.6						
M11	-0.208	-0.188	0.020	64.4	85.0						

Horizon	Target	Quarterly Estimation			
		Mean	Bias	Coverage	
				68%	90%
Q0	-0.082	-0.081	0.001	61.4	84.1
Q1	-0.105	-0.104	0.001	60.5	80.8
Q2	-0.134	-0.134	0.000	58.7	81.4
Q3	-0.190	-0.182	0.008	61.0	82.1

Notes: The first two panels report the baseline (two-shock) employment simulation at monthly and quarterly frequency; the third panel redraws the quarterly simulation with a single shock on the first calendar day of each quarter, perfectly aligned with the aggregation window. The ‘Target’ column displays the average values of the daily Employment IRF displayed in Figure 2 at the corresponding frequency. ‘Mean’ is the average value of the point estimates from local projections fit on aggregated daily data across 1,000 simulations. ‘Coverage’ is the fraction of simulations in which the target falls in confidence intervals of a given size.

F Appendix for Section 7.3

F.1 Mortgagors vs. Non-Mortgagors

This appendix provides supplementary details on the mortgagor versus non-mortgagor analysis presented in Section 7.1. A large literature documents that mortgage debt is a key channel of monetary policy transmission (Cloyne, Ferreira and Surico, 2020; Corsetti, Duarte and Mann, 2022). Motivated by this evidence, we provide descriptive evidence on whether mortgagors and non-mortgagors in our BBVA sample display differential consumption responses to monetary policy shocks. Below we describe the empirical specification and limitations underlying the main-text results.

Empirical Specification. As in the baseline monthly analysis, the dependent variable is the 12-month log change in real consumption (in percentage points), and we include three lags of consumption as controls along with COVID controls and six lags of the monetary policy shock. Following the same approach used for upstream versus downstream sectors in Section 5, we estimate local projections at the customer-month level, interacting the monetary policy shock with an indicator for mortgage status:


$$y_{i,t+h} = \alpha_{h,i} + \sum_{\ell=0}^k \beta_{h,\ell} shock_{t-\ell} + \sum_{\ell=0}^k \gamma_{h,\ell} shock_{t-\ell} \times Mortgage_i + \sum_{\ell=1}^p \varphi_{h,\ell} y_{i,t-\ell} + \theta_h cases_t + \delta_h stringency_t + \varepsilon_{i,t+h}, \quad (20)$$

where $y_{i,t+h}$ is the 12-month log change in consumption for customer i , $\alpha_{h,i}$ denotes customer fixed effects, and $Mortgage_i$ is an indicator equal to one if the customer holds a mortgage. The coefficient $\beta_{h,0}$ captures the response of non-mortgagors; $\gamma_{h,0}$ captures the differential response of mortgagors—our main object of interest. In our preferred specification, we restrict the sample to customers whose average monthly consumption lies between 500 and 15,000 euros. This restriction limits the influence of extreme observations and reduces the likelihood of including accounts that do not represent customers' primary spending accounts. The results are quantitatively very similar when we do not impose this restriction, indicating that the choice of sample bounds does not materially affect our conclusions.

Results. Table 2 in the main text reports the differential consumption response from a specification with customer and time fixed effects. Because time fixed effects absorb all

aggregate variation—including the main effect of the monetary policy shock—that specification identifies only the interaction term and cannot display the individual responses of each group separately. Figure F1 complements the main-text table by estimating the same local projection with customer fixed effects only, so that the responses of mortgagors and non-mortgagors can each be shown. The left panel plots the consumption response of each group; the right panel plots their difference. Both groups display lower consumption in the months following a contractionary shock, but the decline is larger for mortgagors. The differential decreases after the initial months and remains negative—though less precisely estimated—at longer horizons.

Figure F1: Consumption Response to Monetary Policy: Mortgagors vs. Non-Mortgagors




Notes: The left panel plots the consumption response of mortgagors (dashed red) and non-mortgagors (solid blue) to a one standard deviation monetary policy shock, from a specification with customer fixed effects. Lighter-shaded areas are the 90% confidence intervals. The right panel plots the differential response (mortgagors minus non-mortgagors) together with the 68% (lighter-shaded areas) and 90% (darker-shaded areas) confidence intervals of this difference. Confidence intervals are computed using Driscoll-Kraay standard errors. Sample: July 2021–May 2025. The monetary policy shock standard deviation is 2.9bp.

F.2 The Transmission to CPI Disaggregated by COICOP Category

This appendix presents how consumer prices, categorized by COICOP, react to monetary policy changes on a monthly basis. In Figure F2, the impulse responses correspond to these CPI categories: food and non-alcoholic beverages, furnishings, equipment and maintenance, recreation and culture, alcoholic beverages and tobacco, health, education, clothing and footwear, transport, restaurants and hotels, housing and utilities, and communication.

Impulse responses are heterogeneous in economically sensible ways, potentially reflecting differences in the transmission mechanisms and/or the degree of price rigidity across categories. In response to a contractionary monetary policy shock, some categories of price fall markedly. For instance, the CPI for Housing and Utilities and the CPI for Transport experience a pronounced and persistent decline, reaching approximately -0.6% and -0.3%, respectively, by the 11th month. The CPIs for Food and non-alcoholic beverages, Alcoholic beverages and tobacco, Furnishings, equipment and maint., and Restaurants and hotels also decline significantly, albeit their decline is less steep and slightly lagged. In contrast, the CPI for Health, Education and Communication show minimal or no response to the monetary shock.

Figure F2: Monthly response of prices by COICOP category to a monetary policy shock




Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample starts in August 2015 and ends in October 2023. The monetary policy shock standard deviation is 3.7bp.

F.3 Further Evidence on Monetary Transmission Using Monthly Series

This section reports results on the monetary transmission to additional variables relative to the one discussed in the main text. We rely on monthly series focusing on the baseline sample period. Results are shown by Figure F3. The new series include three additional real activity indicators: total new vehicles, net entry of firms, and a retail trade index; three additional expectation variables, concerning unemployment, income, and house prices; and two additional confidence variables, related to the construction and retail sectors. Data on total new vehicles is obtained from EUROSTAT's registrations of new motor vehicles in Spain; both the net entry of firms and the retail trade index are reported by the INE. Expectations regarding unemployment, income, and housing prices are drawn from the ECB's consumer expectations survey, conducted solely at the EA level since April 2020. Lastly, the EU Commission provides the confidence indicators for construction and retail for Spain.

Examining the real activity indicators, we note that each of them sharply decreases in the two months following a monetary policy tightening. New total vehicles and net entry of firms stabilize after this initial decline, notwithstanding some oscillation in point estimates, whereas the retail trade index resumes its decline after the seventh month, which is consistent with our main findings for aggregate daily consumption. Expectations variables as group behave as expected, with households anticipating an increase in unemployment, a decrease in income, and a decline in housing prices. Finally, confidence falls persistently in both the construction and the retail sector.

Figure F3: Monthly response of other selected variables to a monetary policy shock



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. For all series—except for expectations for unemployment, income and house prices which start in April 2020—the sample starts in August 2016 and ends in October 2023. The monetary policy shock standard deviation is 4.2bp for the full sample series and 3.6bp for the two series that start in April 2020.

G Appendix for Section 7.4

The four subsections below report in full the robustness exercises summarized in Section 7.4 of the main text.

G.1 Short-Lags in Industrial Production and Demand

Industrial production is the most commonly available proxy indicator of economic activity at monthly frequency—unsurprisingly, widely employed in the analysis of the transmission of monetary policy. It is a proxy to output, complementing the information from our VAT sales series (particularly for manufacturing and goods producing sectors) while facilitating a comparison of our short-run responses to those reported for other economies. In Figure G1, first panel, we show that, following a monetary shock, monthly industrial production in Spain does contract at short lags, declining by 0.30% on impact and by 0.72% at the 1 month horizon. At longer lags the contraction first moderates, then turns significant again through month seven. Importantly, this pattern is by no means specific to Spanish data; working with US Industrial Production data, Miranda-Agrippino and Ricco (2021) and Bauer and Swanson (2023) report similar short-run responses, consistent with the evidence here.⁷⁵

Using data at monthly frequency, we can also reconsider and extend our headline findings regarding consumption and investment demand. As already noted in Section 2, based on monthly tabulations of VAT sales' declarations by large corporations, the Spanish Tax Authority (i) compiles disaggregated series of sales distinguishing aggregate consumption, aggregate investment and aggregate intermediate input sales and (ii) tabulates external sales, i.e., export and import series.⁷⁶ Thus, at the monthly frequency, we have an alternative data source, complementing our baseline, transaction-based, consumption and investment series.


Results are shown in Figure G1. The sales of intermediate goods and services as well as exports essentially display the same response pattern as the Spanish Industrial Production Index—see the top row of the figure.⁷⁷ The responses of the VAT-derived monthly final consumption measure—as well as of its disaggregated series for consumption goods and consumption services (in the second row of the figure)—are also con-

⁷⁵Appendix C.1 further benchmarks our response against the literature, comparing both its magnitude and shape to other estimates.

⁷⁶These series, sourced from the External Trade in Spain Statistics (AEAT - Spanish Tax Agency) are direct inputs to the compilation of quarterly national accounts by Spain's National Statistical Institute.

⁷⁷The IR of import sales display the same pattern as export. We do not shown the corresponding plot to avoid redundancy.

Figure G1: Monthly IRFs of industrial production and sales



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample starts in August 2015 and ends in October 2023. The monetary policy shock standard deviation is 3.7bp.

sistent with our daily-frequency baseline findings. Final consumption sales drop in the first month of the contractionary monetary policy shock, reaching a local trough in the following month. A second local trough follows later in the year, around 10 months after the shock. The estimated size of the short-lag response is also very similar to our baseline results. The disaggregated series suggest that the brunt of the adjustment is borne out by the consumption of services—broadly in line with our findings in Figure 3, documenting significant short-run adjustment in spending on education, health, recreation and culture, and restaurants and hotels, all of which are consumption services.

The response of the monthly investment series derived from VAT data is shown on

the third and final row of Figure G1. The figure includes the response of aggregate investment (first column), investment in equipment and software (second column), and investment in construction (third column). Relative to our results based on daily data, local projections based on these series are significantly less noisy. Qualitatively, the figure confirms the patterns discussed in Sections 4 and 5.2.⁷⁸ Relative to consumption, the short-lag response of investment is again stronger—primarily driven by spending on equipment, maintenance and software. The muted response in construction investment is consistent with the daily sales response of construction firms in Figure 4.

G.2 Seasonal Adjustment and Smoothing of Daily Series


As discussed in Section 3.2, working with raw daily data presents at least two specific challenges. First, the periodic structure of daily data is less standard, and calendar effects together with (possibly moving) holidays need to be accounted for. Second, daily series are typically much noisier than lower frequency series. This is either because the construction of lower frequency series entails smoothing—through time averaging—of intrinsically high-volatility series; or because measurement error is larger at the daily frequencies. Recall that in our baseline approach we address these two challenges separately: calendar effects are removed via day-of-the-week and day-of-the-month fixed effects (followed by year-on-year growth rates), while daily noise is handled on the estimation side through Smooth Local Projections (see Section 3.2). In this section, we assess the robustness of our baseline results to employing alternative methods to deal with seasonality, calendar effects, and noise.

In a first alternative, we apply a backward-looking 30-day moving average to the raw series—smoothing over daily noise, irregular events and moving holidays—. We then calculate year-on-year growth rates on the resulting series, and estimate standard (unsmoothed) LP via OLS. This approach bundles the treatment of seasonality and noise into a single step, pre-smoothing the data, at the cost of introducing moving-average-induced serial correlation in the dependent variable.

In the second alternative, we adopt an unobserved components approach that econometrically decomposes the raw daily series into distinct components—trend-cycle, seasonal and irregular—with flexible laws of motion for each component; see, for example, Harvey (1989) for a review. By directly estimating seasonal and irregular (i.e. daily noise) components, unobserved-component methods provide a joint solution to seasonality and noise-smoothing challenges: once estimated, subtracting these components from the raw

⁷⁸An immediate contraction, with a short run trough in the second month after the shock, is followed by a spell of stabilization, and then again by a decline towards the end-of-year.

Figure G2: Daily response of real activity to monetary policy shock under alternative seasonal and smoothing procedures



Notes: Responses to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation- consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in October 30th, 2023. The sample begins one year after the start of each series reported in Table A1. For sales and investment (the shorter series) the monetary policy shock standard deviation is 4.1bp, while for consumption and employment it is 3.7bp. The left column reproduces our baseline—day-of-the-week and day-of-the-month seasonal adjustment combined with Smooth Local Projections. The remaining columns report two alternative seasonal adjustment and smoothing methods: a 30-day backward moving average with standard local projections (middle column), and a model-based (TBATS) seasonal adjustment that removes both seasonal and irregular components, with standard local projections (right column).

series yields a seasonally adjusted, de-noised series. In particular, we follow the Spanish Tax Authority’s own suggestion (based on Cuevas, Ledo and Quilis, 2021 treatment of seasonality and noise in daily VAT sales), and implement a variant of the “Trigonometric seasonality, Box-Cox transformation, ARMA innovations, Trend and Seasonality”

(TBATS) unobserved components model proposed in [De Livera, Hyndman and Snyder \(2011\)](#).⁷⁹ We then estimate standard Local Projections via OLS on the filtered series.

In Figure [G2](#), we present the responses of daily sales, consumption, investment, and employment to a monetary policy shock, using our baseline together with two alternative methods. The left column reproduces our baseline—day-of-the-week and day-of-the-month seasonal adjustment combined with Smooth Local Projections—to facilitate comparison. In the middle column, we show results relying on a backward-looking 30-day moving average coupled with standard LP. In the right column (labeled Model-based SA & Smoothing), results are conditional on applying the TBATS unobserved-components methodology, removing both seasonal and irregular components and then estimating standard Local Projections. Reassuringly, the responses of sales, consumption and employment are similar across the three columns and consistent with our baseline findings. Some differences are noticeable in the response of investment, especially for the model-based approach. This most likely reflects the fact that the incidence of noise in the raw data underlying the investment series is stronger relative to the other series. As a result, relative to our baseline, the model-based method tends to interpret short-run investment reactions as noise (while still capturing medium- to long-run reactions).

For completeness, in Appendix [G.5](#), we further extend our analysis and establish that our baseline results are qualitatively and quantitatively robust to a wide range of alternative seasonal and calendar adjustments and alternative smoothing methods.⁸⁰


G.3 Monetary Policy Shocks

In our baseline, monetary policy shocks in the euro area are identified relying on the methodology proposed by [Jarociński and Karadi \(2020\)](#). As discussed in Section [3.1](#), this methodology combines high-frequency identification with sign restrictions in a Bayesian VAR to control for the information effect—its application ensures that the shock series

⁷⁹We proceed in two steps. First, we implement a standard pre-processing step where we remove calendar and other deterministic effects (such as day-of-the-week) via a multiplicative seasonal ARIMA model with automatic detection for outliers and with Spanish public holidays dummies included. Second, we estimate the (stochastic) seasonal and irregular components with TBATS, following closely the specification discussed in [Cuevas, Ledo and Quilis \(2021\)](#): we apply a log-transformation, impose no trend damping, and include weekly, monthly and annual seasonal components with a periodicity of 7, 30.4375 and 365.25 days, respectively. We then remove both the estimated seasonal and irregular components from the original series. Finally, we compute year-on-year growth rates on this filtered series to facilitate comparison of results.

⁸⁰Specifically, in the appendix we systematically explore combinations of alternative seasonal adjustment methods—ranging from no adjustment at all to alternative econometric approaches—with alternative smoothing procedures—ranging from no smoothing at all, to moving averages over different windows, or to exponential smoothing.

Figure G3: Daily response of real activity to alternative monetary policy shocks



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Lighter-shaded areas are the 90% confidence intervals. The sample ends in October 30th, 2023. The sample starts one year after the sample start of each series reported in Table A1. For sales and investment—the shorter series—the baseline, target and poor man’s 1-month OIS monetary policy shock standard deviation is 4.1bp, 3.8bp and 3bp, respectively; while for consumption and employment it is 3.7bp, 3.6bp and 2.9bp, respectively. The 1-month OIS (raw) line corresponds to the corresponding shock series not using the poor man’s approach to control for the information effect.

we use is standard, off-the-shelf and publicly available. In this section, we analyse the robustness of our baseline results to the use of alternative methods to identify monetary disturbances.

We construct two alternative series of monetary policy shocks. In both, we make use of the publicly available Euro Area Monetary Policy Event-Study Database (EA-MPD) provided by [Altavilla et al. \(2019\)](#), recording intraday asset price changes around monetary policy events. The first shock series is the observed 1-month Overnight Indexed Swap (OIS) change around policy decision announcements. In this case, to control for the information channel, we resort to the “poor man’s approach” proposed by [Jarociński and Karadi \(2020\)](#), which consists of excluding observations whenever the sign of the


change in the 1-month OIS is the same as the change in the Euro STOXX50E Index.

The second series of shocks is the Policy Target factor proposed by [Altavilla et al. \(2019\)](#)—which we have extended to cover our baseline sample period. The Policy Target factor refers to the immediate market response to changes in the ECB’s main policy rate, typically announced during the policy decision release. This factor predominantly affects short-term interest rates, causing significant movements at the short end of the yield curve, while having minimal impact on longer-term rates. To construct this alternative series, we follow the methodology proposed in [Altavilla et al. \(2019\)](#), consisting of applying factor analysis to the observed changes in the yield curve.

In Figure [G3](#), we show the responses of sales, consumption, employment and investment to these two alternative monetary policy shock measures. The dynamics implied by the estimated local projections closely track our baseline results, with the responses under alternative shock series contained within the baseline confidence intervals.⁸¹

We conclude noting that, differently from the shocks used in our baseline, the two alternative series in this subsection are arguably better thought of as instruments for monetary policy surprises, rather than actual monetary policy shocks. Indeed, for this type of shock series, [Stock and Watson \(2018\)](#) argue that a LP-Instrumental Variable approach is to be preferred to using instruments as direct measurement of monetary shocks. In Appendix [G.6](#) we report responses obtained from LP-IV, showing that they are similar to those in Figure [G3](#). Altogether, this confirms that our baseline findings on the monetary transmission to daily real activity are robust to alternative methods to identify monetary policy shocks.⁸²

Figure G4: Daily response of real activity before COVID-19



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in December 2019. The sample begins one year after the start of each series reported in Table A1. To make responses comparable to the baseline responses (solid line) reported in Figure 2, the monetary policy shock size is 3.7bp.


G.4 COVID-19

To ensure that our estimates are not disproportionately driven by the extreme observations and idiosyncratic dynamics observed during COVID-19, in our baseline methods we have followed the approach discussed in [Schorfheide and Song \(2024\)](#) and [Lenza and Primiceri \(2022\)](#), prescribing to drop pandemic-period observations. Here we conduct two robustness checks to confirm that our results do not reflect the special features of the COVID-19 subsample.

⁸¹Monetary surprises further out the term structure transmit more slowly: shocks to shorter-maturity rates affect spending more rapidly than shocks to longer-maturity rates, as documented for credit-card spending by [Grigoli and Sandri \(2026\)](#). Our baseline shock isolates the conventional short-rate movements that, operating through the deposit facility rate, drove monetary policy over our sample; we therefore do not pursue a separate analysis across maturities here, and refer the reader to [Grigoli and Sandri \(2026\)](#) for that evidence.

⁸²In an exploratory exercise restricted to the first 90 days after the shock, we also examined asymmetries between contractionary and expansionary surprises: consumption and sales fall significantly after contractionary shocks, while the responses to expansionary shocks are weaker, and the investment and employment responses to expansionary shocks are statistically indistinguishable from zero at the 90% level. These estimates are imprecise: our sample contains few surprises, unevenly split between contractionary (36) and expansionary (27), and the larger surprises are predominantly contractionary; a reliable analysis of asymmetries would require a longer sample. The direction of the asymmetry is consistent with [Grigoli and Sandri \(2026\)](#), who document that spending responds more strongly to contractionary than to expansionary rate surprises in daily credit card data for Germany.

Figure G5: Daily response COICOP consumption categories before COVID-19



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in December 2019. The sample begins one year after the start of each series reported in Table A1. To make responses comparable to the baseline responses (solid line) reported in Figure 3, the monetary policy shock size is 3.7bp.

First, making use of the longest of our daily series—consumption (total and disaggregated by COICOP categories) and employment—we estimate the transmission of monetary policy shocks on the subsample ending in December 2019, before the COVID-19 pandemic. The restriction gives us three and a half years of data: we only have enough variation to estimate responses up to 93 days after a monetary policy shock. At this horizon, local projection estimates are based on 24 monetary policy shocks.⁸³

Figure G4 illustrates our findings. Overall, the short-lags for aggregate consumption and employment in the pre-COVID-19 sample are qualitatively similar to our baseline—only slightly attenuated quantitatively. This suggests that our baseline results are not driven by dynamics specific to COVID-19. This conclusion is corroborated by the re-

⁸³To be clear, extending this exercise to sales and investment is not feasible as the respective samples are significantly shorter.

sponses of disaggregated daily consumption broken down by COICOP categories depicted in Figure G5. Generally, over the first 93 days following a monetary policy shock, the responses by consumption category remain consistently aligned across the two, pre-COVID-19 and full, samples.

In our second robustness check, we rely on the monthly data spanning a much longer, pre-COVID-19 sample period—the data are introduced in Section G.1. We analyze the transmission of monetary policy to monthly total sales, consumption, investment, and employment – our baseline variables – and ten other real economic indicators, over the longest available pre-COVID sample, running from either January 2000 (most variables) or January 2012 (employment variables) to December 2019. Details about the analysis and results are given in Appendix G.7, where we show that the short lag response of key real variables in the longer, monthly pre-COVID-19 sample are qualitatively and quantitatively consistent with our baseline findings.

Overall we conclude that our results on the monetary transmission at short lags is stable across sub-samples, and not driven by dynamics specific to COVID-19 subsample.

G.5 Further Results on Seasonality Adjustment and Smoothing of Daily Series


This appendix complements the analysis of seasonality and smoothing in Appendix G.2. To start with, we explore the sensitivity of our results to specific changes in the smoothing method. Namely, we show results conditional on (i) not smoothing the data; (ii) setting the number of days of the backward MA in our poor man’s approach to 7, 30 (baseline) and 90, respectively; and (iii) applying exponential smoothing. Next, for each smoothing procedure we show results conditional on (i) not seasonally adjusting the data; (ii) applying the TBATS procedure—shown in Appendix G.2; and (iii) estimating a Fractional Airline Model (FAM)—a model-based seasonal adjustment method developed by the Bank of Belgium, as part of an R package with access to a Java library under development for the JDemetra+ program.⁸⁴ We will compare IRFs for all possible combinations of smoothing and seasonal adjustment techniques.

Figures G6, G7, G8, and G9 present our findings under various combinations of seasonal adjustment methods (three, one per column) and smoothing procedures (five, one per row), for sales, consumption, investment, and employment, respectively. To facilitate comparisons, we overlay in every panel, as a dashed black line, our baseline impulse

⁸⁴JDemetra+ is a tool for seasonal adjustment developed by the National Bank of Belgium in cooperation with the Deutsche Bundesbank and Eurostat in accordance with the Guidelines of the European Statistical System (ESS).


responses—those reported throughout the main text, obtained with Smooth Local Projections on day-of-the-week and day-of-the-month seasonally adjusted data. The panel on the third row of the first column corresponds to the approach used in the original submission: a 30-day backward-looking moving average without any direct prior seasonality adjustment. A first takeaway from this comparative analysis is that our baseline impulse responses (the dashed black line) are closely tracked across the large majority of the fifteen seasonal-adjustment and smoothing combinations; in particular, they look remarkably similar to those obtained by applying a 30-day MA smoothing *after* seasonally adjusting the series (compare the graphs on the third row of each figure). This suggests that our baseline approach appears to take care of the most significant seasonal patterns. A second point concerns the role of data smoothing. The first row of our figures shows that, with the exception of employment, seasonal adjustment on its own tends to produce noisy IRFs. Noise in turn tends to reduce the information content of our IRFs—in other words, noise can outweigh seasonality concerns when dealing with high-frequency series of sales, consumption, and investment, motivating smoothing. And yet, one needs to exercise caution in applying smoothing methods to the data. On one hand, smoothing reduces noise; on the other hand, it filters out short-run dynamics. Our graphs on the third row of the figures suggest that the 30-day MA exhibits smooth IRFs while swiftly adapting to changing conditions—both desirable outcomes. In contrast, a 7-day MA (second row of the figures) results in significantly volatile IRFs, while a 90-day MA (fourth row) tends to blur very short-lag responses in all variables—the analysis misses some of the fast responses. Together with the Smooth-LP and model-based (TBATS) results reported in Appendix [G.2](#), this broad similarity across methods confirms that our Smooth-LP baseline is robust to the choice of seasonal adjustment and smoothing procedure.

Figure G6: Daily response to monetary policy shock under alternative seasonal and smoothing procedures: sales




Notes: Responses to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in October 30th, 2023. The sample begins one year after the start of each series reported in Table A1. The monetary policy shock standard deviation is 4.1bp. Across columns, the figure displays the responses under different seasonal adjustment procedures, and across rows, it presents responses under different smoothing methods. The dashed black line in each panel reproduces the paper's baseline—Smooth Local Projections on day-of-the-week and day-of-the-month seasonally adjusted data.

Figure G7: Daily response to monetary policy shock under alternative seasonal and smoothing procedures: consumption




Notes: Responses to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in October 30th, 2023. The sample begins one year after the start of each series reported in Table A1. The monetary policy shock standard is 3.7bp. Across columns, the figure displays the responses under different seasonal adjustment procedures, and across rows, it presents responses under different smoothing methods. The dashed black line in each panel reproduces the paper's baseline—Smooth Local Projections on day-of-the-week and day-of-the-month seasonally adjusted data.

Figure G8: Daily response to monetary policy shock under alternative seasonal and smoothing procedures: investment



Notes: Responses to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in October 30th, 2023. The sample begins one year after the start of each series reported in Table A1. The monetary policy shock standard deviation is 4.1bp. Across columns, the figure displays the responses under different seasonal adjustment procedures, and across rows, it presents responses under different smoothing methods. The dashed black line in each panel reproduces the paper's baseline—Smooth Local Projections on day-of-the-week and day-of-the-month seasonally adjusted data.

Figure G9: Daily response to monetary policy shock under alternative seasonal and smoothing procedures: employment




Notes: Responses to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample ends in October 30th, 2023. The sample begins one year after the start of each series reported in Table A1. The monetary policy shock standard is 3.7bp. Across columns, the figure displays the responses under different seasonal adjustment procedures, and across rows, it presents responses under different smoothing methods. The dashed black line in each panel reproduces the paper's baseline—Smooth Local Projections on day-of-the-week and day-of-the-month seasonally adjusted data.

G.6 LP vs. LP-IV for Instruments for Monetary Policy Shocks

This appendix compares the Impulse Response Functions (IRFs) derived from Local Projections (LP) and LP with Instrumental Variables (LP-IV) for each of the two alternative monetary policy shock instruments found in Figure G3 of the main document. The instruments in question are the target and the 1-month Overnight Index Swap (OIS), adjusted for the information channel via the “poor man’s” approach outlined by Jarociński and Karadi (2020). Unlike our baseline monetary policy shocks, these alternative series serve as instruments for monetary surprises rather than direct measures of monetary policy shocks. According to Stock and Watson (2018), in such situations, simple LP should be replaced by LP-IV methods. In Figure G10, the baseline sample analysis shows no significant difference between LP and LP-IV results over the horizons of up to 250 days after a contractionary monetary policy surprise.

Figure G10: IRFs derived from Local Projections (LP) and LP with Instrumental Variables (LP-IV) for two alternative monetary policy shock instruments



Notes: LP (red) and LP-IV (black) impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Shaded areas are the 68% confidence intervals. The sample starts in August 2015 and ends in October 2023. For sales and investment—the shorter series—the target and poor man's 1-month OIS monetary policy shock standard deviation is 3.8bp and 3bp, respectively; while for consumption and employment it is 3.6bp and 2.9bp, respectively.

G.7 Further Results on Robustness Checks to COVID-19 using Monthly Data

This subsection uses monthly data to provide further robustness checks with respect to COVID-19. We analyze the transmission of monetary policy to total sales, consumption, investment, and employment (our baseline variables) and ten other real economic indicators, over the longest available Pre-COVID sample, running from January 2000 to December 2019. The ten real variables include final consumption sales (disaggregated by goods and services), investment sales (equipment and software, construction), intermediate goods sales, exports, imports, and both permanent and temporary employment, as well as industrial production.

The results for the baseline variables are shown in Figure G11, for the other variables in Figure G12. We find that the responses of the monthly economic indicators are overall significant at short lags, corroborating the robustness of our findings in our baseline sample with daily data. In Figure G11, consumption sales respond on impact, then stabilize and fall again significantly in month 3 after the shock; total sales decline in month 1, investment in month 2. While total employment remains flat, the response of permanent employment (shown below) is smooth, similar to the baseline.

Coming to the disaggregated categories in Figure G12, we highlight the following.

- Final Consumption Sales.

The response pattern of final consumption sales in goods and services mirrors that of total consumption sales (reproduced in Figure for comparison). Both are similar to our Baseline: a significant drop on impact in response to a monetary policy shock is followed by temporary stabilization before again a significant decline from month 3. An important difference between the response of sales of goods and the sales of services, however, is that the contraction in the latter is significant only at the 68% confidence level.

- Investment Sales.

Investment sales overall exhibit a strong response to monetary policy shocks. The response is particularly strong for investment in the construction sector. In the long Pre-COVID sample, the decline in investment in construction is sharp and prolonged following the shock.

- Intermediate Goods Sales, Exports, Imports and Industrial Production.


The response of intermediate good sales and export sales follows a similar pattern, with a significant decline starting in the first month after the shock—somewhat more pronounced for exports. Relative to these two series, imports display a significantly larger negative response, and is already significant on impact at the 68% level. Industrial production also responds in the first month after the shock, with a comparable magnitude and pattern to intermediate goods, export and import sales.

- Employment

The impulse response of permanent employment closely follows the response of total employment in our baseline, with a significant if contained contraction the middle of the year. In the long pre-COVID-19 sample, however, temporary employment do not exhibit the pronounced fall at short lags we detect in our baseline and, most crucially, it exhibits a rise from the middle of the year, if only marginally significant (at the 68% confidence level) in month 12. The pattern in temporary employment may account for the difference in the response of total employment in the long pre-COVID-19 sample, relative to our baseline.


The results from our analysis at monthly frequency extended to a long Pre-COVID-19 sample demonstrates considerable robustness of our results on the short lags of monetary policy.

Figure G11: Monthly response of baseline real activity variables before COVID-19



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample starts in January 2000 and ends in December 2019 and the monetary policy shock size is 3.7bp.

Figure G12: Monthly response of real activity before COVID-19



Notes: LP impulse response functions to a one standard deviation monetary policy shock. The responses are reported in levels. The confidence intervals are computed from heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. Darker-shaded areas are the 68% confidence intervals, and lighter-shaded areas are the 90% confidence intervals. The sample starts in January 2000 and ends in December 2019 and the monetary policy shock size is 3.7bp.