

Household Excess Savings and the Transmission of Monetary Policy*

Thiago R.T. Ferreira, Nils Gornemann, and Julio L. Ortiz
Federal Reserve Board

Household savings rose above trend in many developed countries after the onset of COVID-19. Given its link to aggregate consumption, the presence of these “excess savings” has raised questions about their implications for the transmission of monetary policy. Using a panel of euro-area economies and high-frequency monetary policy shocks, we document that household excess savings dampen the effects of monetary policy on economic activity and inflation, especially during the pandemic period. To rationalize our empirical findings, we build a New Keynesian model in which households use savings to self-insure against countercyclical unemployment and consumption risk.

JEL Codes: E12, E21, E24, E31, E52.

1. Introduction

During the COVID-19 pandemic, households around the world accumulated large stocks of savings through a combination of precautionary motives, an inability to spend their funds amid widespread lockdowns, and increased fiscal support to their incomes. Soon after, the main concern of policymakers quickly switched from alleviating the lack of household income to fighting decades-high inflation. Still,

*We are grateful to Mitch Lott for outstanding research assistance. We also thank Steven Ongena (the editor), Matteo Iacoviello, Phillip Jefferson, Christopher Waller, and presentation attendees at the Federal Reserve Board and the 2024 IJCB conference for their comments and suggestions. The views expressed in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. Author e-mails: thiago.r.teixeiraferrera@frb.gov, nils.m.gornemann@frb.gov, julio.l.ortiz@frb.gov. Address (all authors): Federal Reserve Board, International Finance Division, Washington, DC 20551, USA.

households maintained robust stocks of “excess savings” (i.e., savings in excess of trend), and policymakers began to question their possible effects on the transmission of monetary policy.¹

In this paper, we evaluate the effects of excess savings on monetary policy transmission, both empirically and theoretically. We document that excess savings in euro-area economies rose to historically high levels during the pandemic period. Then, using high-frequency monetary policy shocks, we estimate state-dependent effects of monetary policy using local projections (Jordà 2005), and we find that monetary policy transmission to both inflation and economic activity is dampened in periods of high excess savings. Finally, we rationalize how excess savings affect the transmission of monetary policy using a New Keynesian model in which households face idiosyncratic, countercyclical unemployment and consumption risk against which they can only self-insure through savings.

We begin by measuring the stock of excess savings and monetary policy shocks for euro-area economies. In addition to the euro-area aggregate, our empirical analysis focuses on the four largest euro-area countries: Germany, Italy, France, and Spain. Following de Soyres, Moore, and Ortiz (2023), we define household excess savings as the amount of savings arising from above-trend household savings rates. To estimate country-level trend savings rates, we employ the Hamilton (2018) filter. Our measure of excess savings exhibits variation across time and countries, with the pandemic period showing historically high levels. To measure the monetary policy shocks, we apply the high-frequency approach of Bu, Rogers, and Wu (2021) to the euro area, which accounts for the mix of policies from the European Central Bank (ECB) focused on both policy rates and asset purchases.

After measuring these objects, we use them to estimate the effects of monetary policy on real and nominal outcomes. We focus on two variables of interest: the unemployment rate and consumer price inflation. We estimate the effects of monetary policy on these outcomes of interest via local projections. Our estimates reveal that the effect of a contractionary monetary policy shock on real and

¹For instance, see the speech by Christine Lagarde, President of the European Central Bank, at “The ECB and Its Watchers XXIII” conference.

nominal outcomes of interest is attenuated when stocks of excess savings are larger. We find that our results are robust to (i) different measures of economic activity, (ii) controlling for the balance sheet strength of the banking system, and (iii) time-specific changes around the COVID period.

Our empirical results can be summarized as follows. We find that a contractionary monetary policy shock scaled to generate a 50 basis point increase in the two-year German government yield raises the unemployment rate by about 0.30 percentage point when excess savings are close to zero, but only by about 0.15 percentage point when excess savings are fixed to their 2023:Q1 levels. Additionally, we show that 12-month headline inflation declines by nearly 0.40 percentage point when excess savings are close to zero, but only by about 0.30 percentage point when excess savings are consistent with their observed 2023:Q1 levels. We choose to evaluate the efficacy of monetary policy in 2023:Q1 because euro-area headline inflation peaked around this time, thereby constituting a moment at which policymakers paid heightened attention to the effectiveness of monetary policy going forward.

Motivated by our empirical findings, we build a simple New Keynesian model with unemployment and imperfect insurance against individual unemployment risk. Relative to the standard representative agent New Keynesian setting, savings in our model are valued because they allow workers to self-insure against the consumption risk of being unemployed. Unemployment risk rises during contractions, inducing workers to cut back consumption further at the start of a downturn to save more. This response amplifies the direct impact of any contractionary shock. Higher savings at the onset of the downturn reduce this consumption response and therefore lead to less amplification. As a result, consistent with our empirical findings, a calibrated version of our model generates dampened real and nominal responses to monetary policy shocks in a high-savings economy relative to the baseline economy. In total, our empirical and quantitative results point to an economically meaningful nonlinearity in the potency of monetary policy based on the level of stocks of excess savings.

Our paper relates to the recent literature measuring excess savings and studying their aggregate implications. Our method for measuring excess savings follows de Soyres, Moore, and Ortiz

(2023). Alternative approaches to measuring excess savings include Aladangady et al. (2022) and Abdelrahman and Oliveira (2023). We take a time-varying filtering approach because it best suits our purpose of estimating the effects of monetary policy shocks in a panel of euro-area economies with a rich time dimension. Additionally, there are other recent studies evaluating the effects of excess savings accumulated during COVID-19. Auclert, Rognlie, and Straub (2023) analyze the process by which household excess savings affect the level of aggregate demand, and how this process varies based on the distribution of excess savings. Aggarwal et al. (2023) study debt-financed fiscal transfers in a model of the world economy that reproduces large fiscal deficits, large increases in private savings, and persistent current account deficits. We contribute to this literature by empirically and quantitatively linking aggregate excess savings to monetary policy effectiveness in the context of the euro area.

Our paper also relates to the literature on how monetary policy transmits to households. Recent contributions, such as those made by Cloyne, Ferreira, and Surico (2020), find that mortgagors are more sensitive to monetary policy than outright owners because the former have little liquid wealth. In addition, Alpanda, Granziera, and Zubairy (2021) and Harding and Klein (2022) find empirically that monetary policy is more effective in affecting the macroeconomy when household debt is rising or high.² On the modeling side, our paper relates to household models of precautionary savings demand in the presence of countercyclical idiosyncratic risk, such as Gornemann, Kuester, and Nakajima (2016), Challe et al. (2017), Ravn and Sterk (2017), Bilbiie (2018), Den Haan, Rendahl, and Riegler (2018), Acharya and Dogra (2020), and Cho (2023). The empirical state dependence we document can be viewed as supportive of the mechanisms in these papers. Our paper is similar to them, though it studies state-dependent household responsiveness to monetary policy across major euro-area countries and in the context of

²Harding and Klein (2022) provide an alternative mechanism to rationalize our findings. In their model, higher savings would relax a collateral constraint, leading to a dampening of contractionary shocks when constrained agents cut back on consumption. We view our mechanism as complementary to theirs.

excess savings rather than household debt or net worth. Our measure of excess savings, particularly around COVID-19, likely reflects an influx of liquid savings that allowed many European households to remain clear of borrowing constraints during the recent tightening cycle.

The rest of the paper is organized as follows. Section 2 discusses the data used for our empirical analysis—particularly our measures of excess savings and euro-area monetary policy shocks. Section 3 reports our empirical results, which show that the effects of monetary policy are dampened when stocks of household excess savings are high. Section 4 explores the economic mechanism through which excess savings affect the transmission of monetary policy in a simple New Keynesian model and presents simulations rationalizing our empirical results. Section 5 concludes.

2. Data Description

In this section, we describe how we measure the stocks of excess savings and monetary policy shocks used for our empirical results. Appendix A provides further details on the data used in our analysis.

2.1 *Excess Savings in the Euro Area*

We follow de Soyres, Moore, and Ortiz (2023) by defining the stock of excess savings as the amount of assets, as a percent of gross domestic product (GDP), arising from above-trend savings rates. First, for each country, we extract a trend of the savings rate using the Hamilton (2018) filter.³ We then use the detrended savings rate to calculate the flow of excess savings for country i in quarter t , in euros, as follows:

$$\begin{aligned} \text{Flow of excess savings}_{it} &= (\text{Detrended Savings Rate}_{it}) \\ &\times (\text{Disposable Income}_{it}). \end{aligned} \quad (1)$$

³Following Hamilton (2018), we detrend the savings rate country by country via a regression of the savings rate on lags 8 to 11. The residual is the detrended savings rate.

To construct the measure of the stock of excess savings, we calculate the cumulative sum of the flows defined in Equation (1) and normalize it by nominal GDP:

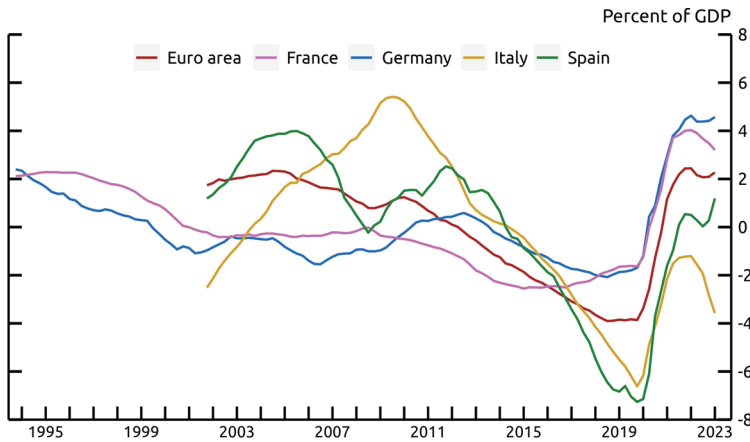
$$\text{Stock of excess savings}_{it} = \frac{\sum_{t=1}^T \text{Flow of excess savings}_{it}}{\text{Nominal GDP}_{it}} \times 100. \quad (2)$$

Finally, we demean the series at the country level over the entire sample period, which spans from 1999:Q1 through 2023:Q2.⁴ By construction, the stock of excess savings increases when the flow of excess savings is positive, while the stock of excess savings decreases when the flow of excess savings is negative.

This measure of aggregate excess savings has some advantages. First, its construction requires only aggregate nominal household savings, disposable income, and nominal GDP, all of which are readily available for a variety of euro-area countries. Second, our methodology produces a full time series of estimated excess savings, allowing us to exploit its variation over time in our analysis. Finally, despite using nominal household savings as an input, our measure does not rely on the assumption that prices remain at their trends.

Our measure of the stock of excess savings for the euro area exhibits considerable variation, both in the time series and across countries. Figure 1 shows that in the lead-up to the 2008–09 Global Financial Crisis, different economies had different trajectories, with Italy increasing its excess savings, France and Germany maintaining their excess savings, and Spain and the euro-area aggregate decreasing their savings. In contrast, from 2012 to 2020, most of these economies ran down their excess savings, with the exception of France. Finally, amid the COVID-19 pandemic, all of these economies saw their stocks of excess savings sharply increase.

⁴We demean the stock of excess savings within country to ensure that our empirical results, which we describe in Section 3, are not driven by permanent heterogeneity in excess savings stemming from initial conditions when accumulating the savings flows over time.

Figure 1. Stock of Excess Savings in the Euro Area

Note: Figure 1 shows the series of stocks of excess savings for the euro-area aggregate (red), France (pink), Germany (blue), Italy (yellow), and Spain (green).

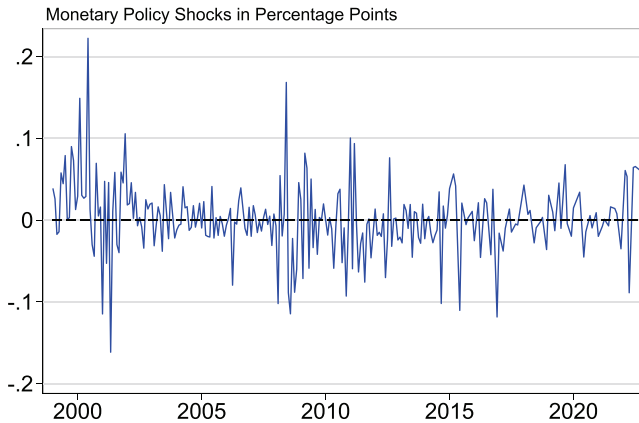
2.2 Euro-Area High-Frequency Monetary Policy Shocks

We apply the methodology of Bu, Rogers, and Wu (2021) to calculate high-frequency monetary policy shocks for the euro area. Specifically, we use daily data from the German Treasury yield curve for the monetary policy meetings of the ECB for the period of January 1999 through September 2022 in our application. Figure 2 displays the monetary policy shocks that are calculated using this methodology. For our regressions, we aggregate these shocks by summing them to a monthly frequency.⁵

Our chosen approach to measuring monetary policy shocks combines three important features that address issues extensively discussed in the literature. First, the shocks bridge periods of conventional and unconventional monetary policy by using interest rate movements from the entire Treasury yield curve. Second, Bu, Rogers, and Wu (2021) provide evidence that this methodology removes the central bank information effect for the U.S.: monetary policy announcements may reveal information about the state of the macroeconomy, instead of representing only genuine monetary

⁵When analyzing the consumption response to monetary policy in Appendix B.2, we aggregate the monetary policy shocks to a quarterly frequency.

Figure 2. Monetary Policy Shocks from the European Central Bank



Note: Figure 2 shows the time series of monetary policy shocks calculated using German Treasury yield curve for the monetary policy meetings of the ECB for the period of January 1999 through September 2022.

policy surprises. Third, Bu, Rogers, and Wu (2021) document that their U.S. monetary policy shocks are not predicted by available information on the economy, such as Blue Chip forecasts, news releases, and consumer sentiment.⁶ Following the same approach, in Appendix B.1 we provide evidence that our euro-area monetary policy shocks are not predicted by information available in real time.

3. Empirical Results

In this section, we use our measures of excess savings and euro-area monetary policy shocks to document that excess savings dampen the transmission of monetary policy to both economic activity and inflation.

3.1 Regression Specifications

Our sample runs from January 1999 through September 2022 and covers five economies: the euro-area aggregate, Germany, France,

⁶For more discussion of these issues, see Miranda-Agrippino (2016) and Bauer and Swanson (2020).

Italy, and Spain. We represent an economy with i and a given month with t . The euro-area monetary policy shock is denoted by ε_t^m and is scaled such that the shock generates a 50 basis point increase in the two-year German government bond yield. For all regressions, we also specify a set of country-specific and euro-area-specific controls denoted by \mathbf{Z}_{it} . The country-specific controls are 12 lags of inflation rates, unemployment rates, industrial production, GDP growth, and an interaction of GDP growth with the monetary policy shock to account for state-dependent effects of monetary policy on economic activity (Tenreyro and Thwaites 2016).⁷ The euro-area controls include 12 lags of inflation, the unemployment rate, and the spread between the five-year triple-B-rated bond yield and the five-year German government bond yield.

We estimate the transmission of monetary policy to measures of economic activity and inflation unconditionally and conditional on the ex ante stock of excess savings. Specifically, we estimate the following local projections (Jordà 2005) at a monthly frequency for a series of horizons h :

$$Y_{it+h} - Y_{it-1} = \beta_i^h + \beta_1^h \varepsilon_t^m + \beta_2^h (\text{Excess Savings Stock}_{it-1}) \times \varepsilon_t^m + \gamma^h \mathbf{Z}_{it-1} + e_{it+h}, \quad (3)$$

where Y_{it} is the measure of economic activity and inflation.

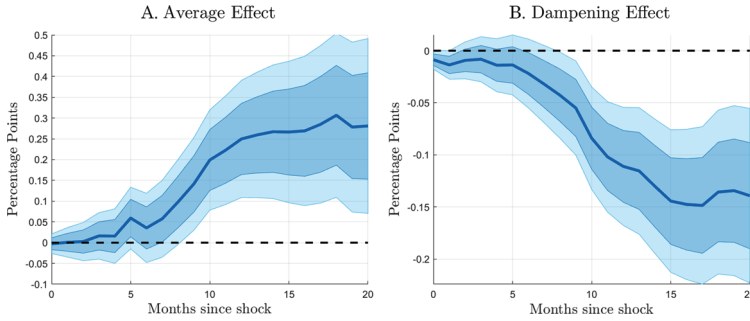
While we focus on the unemployment rate as a measure of economic activity, Appendix B documents that our results are robust to using both aggregate consumption and industrial production instead. In addition, we also show that our results are robust to controlling for the COVID period as well as bank balance sheet strength, respectively.

3.2 *Excess Savings Dampen Monetary Policy Effects on Activity*

Using the unemployment rate as a measure of economic activity, we find that euro-area monetary policy shocks increase the unemployment rate but less so when excess savings are high. Figure 3A plots

⁷We convert GDP from a quarterly frequency to a monthly frequency by assigning the value of realized GDP in a given quarter to every month of that quarter.

Figure 3. Effect of a Contractionary Monetary Policy Shock on Euro-Area Unemployment Rate



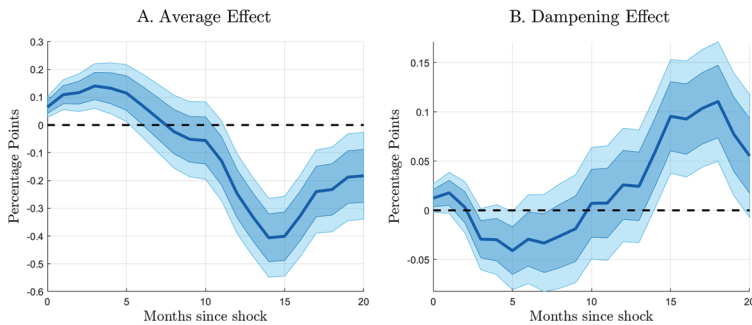
Note: Figure 3 depicts the response of the unemployment rate to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect, β_1^u , from local projections (3), while panel B plots the effect conditional on the level of excess savings, β_2^h . The shaded area reflects the 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay (1998).

the unconditional effects of monetary policy shocks, with the unemployment rate rising over time, reaching its peak effect after 15 to 20 months, and increasing by 30 basis points. Figure 3B plots the estimate of β_2^h , which captures the nonlinearity in monetary transmission with respect to excess savings. More precisely, when the stock of excess savings of a euro-area economy increases by 1 percentage point of GDP relative to the historical average, we find that the effect of the monetary policy shock on the unemployment rate is dampened by roughly one-half. Overall, our estimated effects of monetary policy shocks on economic activity are comparable in magnitude to those from other papers (e.g., Badinger and Schiman 2023).

3.3 *Excess Savings Dampen Monetary Policy Effects on Inflation*

In addition to dampening the response of real outcomes, we find evidence that excess savings also dampen the response of prices. We estimate Equation (3) for 12-month headline inflation and plot the results in Figure 4. The unconditional effect (Figure 4A) shows that inflation declines by 40 basis points in response to a contractionary

Figure 4. Effect of a Contractionary Monetary Policy Shock on Euro-Area Inflation



Note: Figure 4 plots the response of inflation to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect, β_1^h , from local projections (3), while panel B plots the effect conditional on the level of excess savings, β_2^h . The shaded area reflects the 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay (1998).

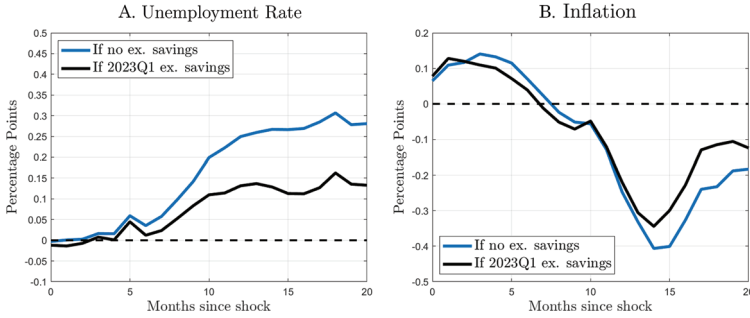
monetary policy shock. When the stock of excess savings of a euro-area economy is 1 percentage point of GDP relative to its historical average (Figure 4B), the decline in inflation is dampened by around 10 basis points. Our results are consistent with the literature documenting that using a high-frequency measurement of shocks helps to show that inflation *decreases* after contractionary monetary policy shocks (e.g., Ramey 2016, Jarociński and Karadi 2020).

3.4 Monetary Policy Transmission Post-Pandemic

We next quantify the effect of the rise in excess savings on the transmission of monetary policy during the COVID-19 recovery. With 12-month headline inflation in the euro-area aggregate having peaked in 2022:Q4 and excess savings still remaining elevated in the same period, we use our estimates to quantify the effectiveness of monetary policy from the perspective of a policymaker in 2023:Q1 as they start to evaluate how tight policy would need to be going forward.

Figure 5 shows that excess savings dampened the effects of monetary policy during the recovery from the pandemic. As a baseline, the blue lines depict the average response of the unemployment and

Figure 5. Effect of a Contractionary Monetary Policy Shock in 2023:Q1



Note: Figure 5 plots the response of the unemployment rate and inflation to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panels A and B plot the total monetary policy effect under two scenarios: (i) when excess savings are equal to zero (i.e., β_1^h from local projections (3)) and (ii) when excess savings are set equal to their 2023:Q1 level based on an average of the countries in our sample.

inflation rates to a monetary policy shock under the assumption that excess savings are at their historical averages (the same as in Figures 3A and 4A, respectively). The solid black lines reflect the response of the unemployment and inflation rates to a monetary policy shock under the assumption that excess savings are at their 2023:Q1 levels. Comparing the two sets of responses, we estimate peak-dampening effects of about one-fourth to one-half on the efficacy of monetary policy for both the unemployment and inflation rates.

4. Excess Savings in a Simple New Keynesian Model

In this section, we build a model that provides an interpretation of our empirical results: higher excess savings flatten the IS curve. Our model is a simple New Keynesian model with equilibrium unemployment due to matching frictions in the labor market. Workers face idiosyncratic, countercyclical unemployment risk against which they can only be partially insured by saving while employed. When unemployment risk rises, households want to cut consumption today to save more, which then amplifies the response of the economy to

the initial shock. However, higher savings dampen this amplification by allowing for better risk sharing. Our exact setup is a simplified version of the model in Challe et al. (2017) and is similar to, for example, Ravn and Sterk (2017) and Heathcote and Perri (2018). The same forces can also be found in less stylized models of precautionary savings over the business cycle like those of Gornemann, Kuester, and Nakajima (2016), Den Haan, Rendahl, and Riegler (2018), and Cho (2023).

4.1 *Timeline*

The economy is populated by a unit mass of identical families. Each family itself consists of a unit mass of workers. At the beginning of the period, the family redistributes bonds between its members who were employed last period. Unemployed members of the family also might hold some bonds or debt, depending on their history. Importantly, they are not able to share in the wealth of the family until they find a job.⁸ After the redistribution of bonds takes place, aggregate shocks realize and firms announce how many workers they plan to hire. Employed workers produce for the firms they work for and are paid. Unemployed workers receive unemployment benefits. Both groups then decide how to split their income between consumption and savings. At the end of the period all employed members of a family meet and pool resources, which are then reshuffled between these family members for the next period.

4.2 *Representative Family*

We describe the problem of the representative family in two stages. The individual state variable of a family's problem is a distribution μ of workers over (N, b) , where N indicates the employment status of a worker and b the current bond holdings of the worker. $N = 0$

⁸The assumption of a representative family that can share its resources within a subset of workers is clearly unrealistic. It is meant to capture a world in which workers self-insure through saving against unemployment risk. The assumption that the employed workers are able to pool resources, however, strongly simplifies the analysis. This simplification arises because a worker's asset position only depends on their current employment status and the length of his current unemployment spell—not, for example, on the length of an employment spell or other unemployment spells in the past.

denotes a worker who is currently employed, while $N > 0$ lists the number of periods a worker has been unemployed. b takes values in $[-\bar{b}, \infty)$, with \bar{b} being the borrowing limit.

At the beginning of the period the value function of the family is $\tilde{V}_t(\tilde{\mu})$ and evolves to the value function $V_t(\mu)$ as labor market flows occur. The household makes no decisions at this stage. Employed workers lose their match with their employment agency with probability λ but are immediately allowed to search for a new one. The probability of finding a job is f_t for both newly and previously unemployed workers and is determined in equilibrium as described below. The resulting transition and relation between value functions is given by the following set of equations:

$$\begin{aligned} \tilde{V}_t(\tilde{\mu}) &= V_t(\mu) \\ \text{subject to } \mu(b, 0) &= (1 - \lambda(1 - f_t))\tilde{\mu}(b, 0) + \sum_{i=1}^{\infty} f_t \tilde{\mu}(b, i) \\ \mu(b, 1) &= \lambda(1 - f_t)\tilde{\mu}(b, 0) \\ \mu(b, N) &= (1 - f_t)\tilde{\mu}(b, N - 1), \text{ for } N > 1. \end{aligned}$$

In the second step, after labor market transitions have taken place, we reach the production and consumption stage. Currently employed workers receive wages w_t , while unemployed workers receive unemployment benefits χ . All pay a proportional tax τ_t on these incomes to finance unemployment benefits and lump-sum taxes T_t to pay for government debt. Finally, they either receive interest income (if b is positive) or repay their debt (if b is negative). These payments are $\frac{R_{t-1}}{\pi_t}b$, where π is the inflation rate and R_{t-1} the nominal interest rate determined last period. Given these incomes, the family assigns a consumption ($c(b, N)$) and savings ($b'(b, 0)$) plan for each (b, N) . These plans have to be consistent with individual budget sets. Unemployed workers carry their remaining savings to the next period, while employed workers meet at the end of the period and pool resources such that all employed workers finish the period with the same share of total savings as the employed.⁹

⁹Concavity in the utility from consumption and the identical probability of losing a job make it optimal for the planner to assign every employed worker

The resulting optimization problem is given by the following equations:

$$\begin{aligned}
 V_t(\mu) = & \max_{(b',N),c(b,N)}_{(b,N) \in \text{sup}(\mu)} \left[\sum_{(b,N) \in \text{sup}(\mu)} \frac{c(b,N)^{1-\sigma}}{1-\sigma} \right. \\
 & \left. + \beta \mathbb{E}_{t+1|t} \tilde{V}_{t+1}(\tilde{\mu}) \right] \\
 \text{subject to } & c(b,0) + b'(b,0) = w_t(1 - \tau_t) + b \frac{R_{t-1}}{\pi_t} - T_t + \Pi_t \\
 & c(b,N) + b'(b,N) = \chi(1 - \tau_t) \\
 & \quad + b \frac{R_{t-1}}{\pi_t} - T_t + \Pi_t, \text{ for } N > 0 \\
 & b'(b,N) \geq -\bar{b} \\
 & \tilde{\mu}(b,0) = \int_{\bar{b}} \mu(\tilde{b},0) d\tilde{b}, \text{ if } b = \frac{\int_{\bar{b}} b'(\tilde{b},0) \mu(\tilde{b},0) d\tilde{b}}{\int \mu(\tilde{b},0) d\tilde{b}} \\
 & \tilde{\mu}(b,0) = 0, \text{ if } b \neq \frac{\int_{\bar{b}} b'(\tilde{b},0) \mu(\tilde{b},0) d\tilde{b}}{\int \mu(\tilde{b},0) d\tilde{b}} \\
 & \tilde{\mu}(b,N) = \mu(b'^{-1}(b,N), N), \text{ if } N > 0.
 \end{aligned}$$

Looking over the family's problem, the pooling of resources is contained in the transition equation for $\tilde{\mu}$, with all employed workers concentrating in one $(b,0)$ pair, while all other $(b,0)$ pairs have zero mass. The last line captures that unemployed households carry their remaining savings forward. These assumptions, together with a tight-enough borrowing constraint, are what make our framework very tractable, as we do not have to follow a large asset distribution or the full employment history of workers as separate state variables. Bonds are valued not only for their interest income but also for their ability to provide consumption insurance to workers who

the same share and to give identical plans to workers with the same (b,N) consistent with optimality, so these assumptions are only made the streamline the presentation.

become unemployed and are temporarily unable to pool resources with the family. As they run down their savings, they eventually hit the borrowing constraint. At that point it does not matter anymore for optimal behavior how long a worker has been unemployed. On the flip side, once a worker becomes employed, they end the period with the same number of bonds as all the other employed workers, making it unnecessary to track their history or the different bond amounts for employed workers. As a result, we only need to follow the mass of employed workers and a finite set of unemployment duration and savings pairs.¹⁰

4.3 Bond Supply

The government issues a constant number of nominal one-period bonds, \bar{B} , each period, which it sells to households and finances through lump-sum taxes on all workers:

$$T_t = \left(\frac{R_{t-1}}{\pi_t} - 1 \right) \bar{B}.$$

4.4 Labor Market Model

This subsection describes the labor market in our model.

4.4.1 Employment Agencies

Employment agencies hire workers by posting vacancies, which are filled at rate q_t . An agency that is matched with a worker rents them out to intermediate goods producers and receive a compensation in the amount of h_t in exchange. It pays the worker a wage w_t while they are matched. The match continues into the next period with probability $(1-\lambda)$. When the match is dissolved, the worker becomes unemployed. We assume that all firms discount future payment flows

¹⁰Our results should carry over to models without risk sharing in a representative family, where instead employed workers keep accumulating savings on their own for insurance against unemployment as in Gornemann, Kuester, and Nakajima (2016).

at the ex ante real rate (r_t). The value to the agency of an ongoing match is J_t and given recursively by the following expression:

$$J_t = (h_t - w_t) + (1 - \lambda) \frac{1}{r_t} \mathbb{E}_{t+1|t} J_{t+1}.$$

Assuming free entry for new agencies, employment agencies keep posting new vacancies until the cost of opening a vacancy, κ , equals the chance of matching times the value of a match:¹¹

$$\kappa = q_t J_t.$$

Finally, we assume that wages follow a simple wage rule, which sees wages rise in total employment (N_t):

$$\log(w_t) - \log(\bar{w}) = \phi_w (\log(N_t) - \log(\bar{N})).$$

In models with matching frictions, many wage determination rules are consistent with equilibrium. We follow Gornemann, Kuester, and Nakajima (2016) and Challe et al. (2017) in picking a parsimonious formulation.

4.4.2 Labor Market Flows

Next, we describe the aggregate labor market flows. We assume that the total number of matches follows a standard matching function:

$$M_t = \mu_M v_t^{\alpha_M} (\lambda N_{t-1} + (1 - N_{t-1}))^{1-\alpha_M}.$$

It takes the number of posted vacancies and the mass of workers searching for employment as inputs.¹²

As a result, the chance of a worker finding a job is

$$f_t = \frac{M_t}{\lambda N_{t-1} + (1 - N_{t-1})},$$

¹¹We assume here implicitly that employment agencies always have a large-enough present discounted value to want to post some vacancies, which will be the case in all our experiments. This also implies that the agencies would never want to end a match with a worker endogenously.

¹²Technically, we should write the matching function as $M_t = \max\{v_t^{\alpha_M} (\lambda N_{t-1} + (1 - N_{t-1}))^{1-\alpha_M}, v_t, (\lambda N_{t-1} + (1 - N_{t-1}))\}$ to rule out cases in which more matches than posted vacancies or searching workers are generated. However, given our calibration, these cases never occur in our experiments.

while the chance of an employment agency finding a worker is

$$q_t = \frac{M_t}{v_t}.$$

Total employment, N_t , evolves as follows:

$$N_t = (1 - \lambda)N_{t-1} + M_t.$$

4.5 Production

Production has two stages. Final goods producers aggregate intermediate goods into a final good that can be used for consumption, vacancies, and price adjustment costs. Intermediate goods producers each create a variety of the intermediate good using labor services as their sole input. They are monopolists for the sale of their variety and are subject to price adjustment costs, generating a source of price rigidity.

4.5.1 Final Goods Producers

Final goods producers sell output Y_t at price P_t produced from a continuum of intermediate varieties in quantities $y_{i,t}$ bought at prices $p_{i,t}$. They solve the maximization problem:

$$\begin{aligned} \max_{Y_t, (y_{i,t})_{i=0}^1} \quad & P_t Y_t - \int_0^1 p_{i,t} y_{i,t} di \\ \text{subject to } Y_t = \quad & \left(\int_0^1 y_{i,t}^{\frac{\nu-1}{\nu}} di \right)^{\frac{\nu}{\nu-1}}. \end{aligned}$$

4.5.2 Intermediate Goods Producers

Intermediate goods producers are producing their variety using a linear technology with productivity \bar{Z} that takes labor services ($n_{i,t}$) as inputs at a price h_t paid to employment agencies. They are monopolists for their variety, setting their price $p_{i,t}$ subject to price adjustment costs while taking final goods producers' demand response into account. As this makes their price a state variable, they optimize the intertemporal value of their profits discounted at the ex ante real

rate (r_t). We denote this value by $J_{i,P,t}$ resulting in the following optimization problem in period t :

$$J_{i,P,t}(p_{i,t-1}) = \max_{y_{i,t}, n_{i,t}, p_{i,t}} \left[p_{i,t} y_{i,t} - P_t h_t n_{i,t} - P_t \Phi_\pi \left(\frac{p_{i,t}}{p_{i,t-1}} - \bar{\pi} \right)^2 + \frac{1}{r_t} \mathbb{E}_t J_{i,P,t+1}(p_{i,t}) \right]$$

subject to $y_{i,t} = \bar{Z} n_{i,t}$

$$y_{i,t} = \left(\frac{p_{i,t}}{P_t} \right)^{-\nu} Y_t.$$

4.6 Monetary Policy

Monetary policy sets the nominal interest rate based on a standard inertial Taylor rule subject to an i.i.d. normal monetary policy shock (ϵ_t^R):

$$\log(R_t) - \log(\bar{R}) = \phi_R (\log(R_{t-1}) - \log(\bar{R})) + (1 - \phi_R) [\phi_\pi (\log(\pi_t) - \log(\bar{\pi}))] + \epsilon_t^R.$$

4.7 Market Clearing and Consistency

Final goods markets clear:

$$Y_t = C_t + \kappa v_t + \Phi_\pi \left(\frac{P_t}{P_{t-1}} - \bar{\pi} \right)^2.$$

Labor services markets clear:

$$N_t = \int_0^1 n_{i,t} di.$$

Aggregate employment is consistent with the distribution:

$$N_t = \mu(\tilde{b}, 0) d\tilde{b}.$$

Bond markets clear:

$$\bar{B} = \int b'_t(b, N) d\mu_t(b, N).$$

Profits are consistent:

$$\Pi_t = Y_t - \kappa v_t - \Phi_\pi \left(\frac{P_t}{P_{t-1}} - \bar{\pi} \right)^2 - w_t N_t.$$

Aggregate and individual consumption are consistent:

$$C_t = \int c_t(b, N) d\mu_t(b, N).$$

Unemployment benefits are paid period by period through τ_t :

$$\tau_t = \frac{\chi(1 - N_t)}{\chi(1 - N_t) + w_t N_t}.$$

4.8 *Some Intuition*

We solve the model using linearization. To keep things really simple, we chose a calibration in which it is optimal for unemployed workers (close to steady state) to deplete all their savings in one period such that, effectively, there are only two relevant Euler equations—the one for workers who have been employed for multiple periods and the one for workers who just found a job—while unemployed workers essentially behave in a hand-to-mouth way. As a result, we only need to track total employment in the economy.

The key departure from a standard representative agent model is that, in our setting, savings are valued, as they allow the family to better insure its members against the consumption risk from being unemployed. The presence of this risk alters the Euler equation(s) of the model. The trade-off characterizing the choice of savings by an employed household at the margin is given by the following:

$$c_t(b, 0)^{-\sigma} = \beta \mathbb{E}_{t+1|t} \frac{R_t}{\pi_{t+1}} \left[(1 - \lambda(1 - f_{t+1})) c_{t+1}(\hat{b}_t, 0)^{-\sigma} + \lambda(1 - f_{t+1}) c_{t+1}(\hat{b}_t, 1)^{-\sigma} \right].$$

Assume that overall savings are low enough that $c_{t+1}(\hat{b}_t, 1) < c_{t+1}(\hat{b}_t, 0)$ —i.e., that consumption of the newly unemployed worker is lower than the worker who remains employed. Then, everything else being the same, a fall in the expected job-finding rate increases the value of the right-hand side of the equation. As a result, today's

consumption has to fall to allow the equation to continue to hold, putting downward pressure on aggregate demand and amplifying the original shock. The more savings households have, the smaller the gap in consumption will be between the two labor market states—and therefore the smaller the push from a fall in expected job-finding rates on today's consumption and, as a result, the lesser the amplification. Thus, higher (excess) savings lead to a weaker output response to any shock and, through the Phillips curve implied by our intermediate goods producers, a weaker inflation response. Now, in the model, wages, taxes, profits, and interest rates will also change, which could possibly dampen the amplification we just described. Therefore, to gauge if we actually generate a quantitative difference in magnitudes, we calibrate and simulate our model for different savings levels.

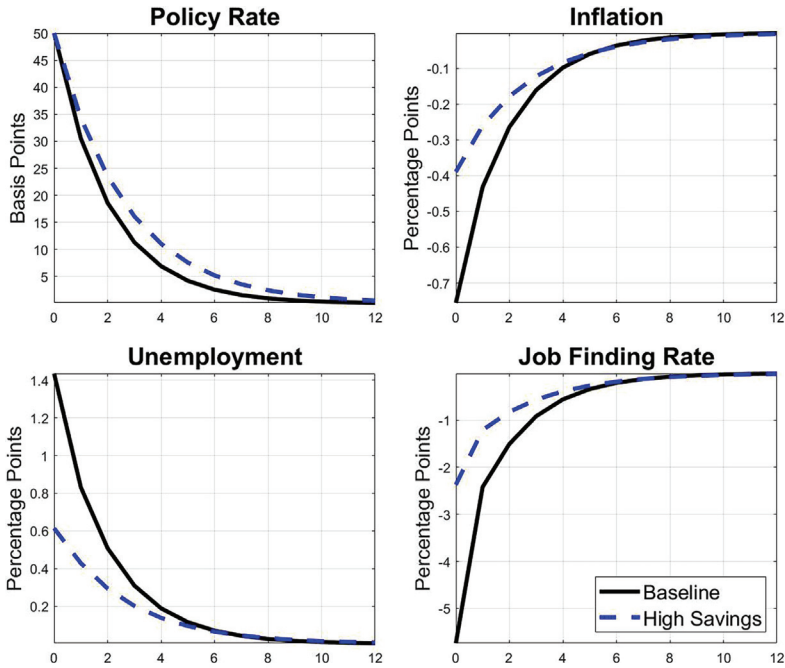
4.9 Calibration and Simulation Results

A period in the model is a quarter. We pick parameters with typical targets in the literature in mind and calibrate the steady state around which we linearize. We target a real annualized rate of 2 percent and an inflation rate of 2 percent. Unemployment is 5 percent and we assume a job-finding rate of 80 percent. We set λ and the scale of the matching function to achieve the latter target in concert with our other parameters. We assume a curvature of the matching function of $\alpha_M = 0.5$. For simplicity we set the borrowing constraint equal to zero and set government bonds equal to 5 percent of output.¹³ Risk aversion σ is 1 and we choose the household's β to be consistent with our real rate target. We normalize \bar{Z} so that output is 1 in the steady state and choose $\phi_w = 0.5$ to have some real wage rigidity. We set $\nu = 3$, a high value in the New Keynesian literature, but one that allows us to obtain a plausible labor share while also having a low ratio of job posting cost to GDP, in line with the literature.¹⁴ We target a labor share of 66 percent given

¹³The latter choice was made for numerical convenience. It is meant to capture the low liquid savings levels of the median household, not the overall stock of government debt.

¹⁴This value is at the lower end of the estimates in the literature and, in our simple setting, has to absorb parts of the absence of capital, investment, and the return to it.

Figure 6. Model Response to a Contractionary 50 Basis Point Monetary Policy Shock



Note: The black line denotes the response under our baseline calibration, while the blue line shows the response when savings as a share of annual GDP are 5 percent higher. Policy rate and inflation responses are annualized.

the other targets in steady state, and χ is set equal to 50 percent of the resulting steady-state wages. We set Φ_π such that the slope of the linearized Phillips curve is the same as in a Calvo model with prices lasting, on average, for a year. Finally, we choose $\phi_R = 0.8$ and $\phi_\pi = 1.1$.¹⁵

The impulse responses in Figure 6 show the impulse response to a monetary policy shock that increases the policy rate on impact

¹⁵1.1 is a low response to inflation in a Taylor rule. It has the advantage that the interest rate paths after a monetary policy shock are roughly the same under different levels of savings, making the comparisons more straightforward. The results would be similar if we allowed for a stronger response but adjusted shocks to generate similar realized paths of the nominal rate.

by 50 basis points (annualized). Inflation and the job-finding rate fall, while unemployment rises on impact. All variables then converge back to normal after roughly two years. The red line shows the results if we recalibrate our model to have 5 percent higher bonds relative to annual GDP in the steady state.¹⁶ As we can see, the impact responses of inflation, unemployment, and the job finding rate are roughly halved, demonstrating that consumption insurance against unemployment risk provides a plausible interpretation of our results. In a representative agent model, in which Ricardian equivalence would hold and idiosyncratic risk would be fully insured, the increase in savings would be instead inconsequential.¹⁷

5. Conclusion

Monetary policy effectiveness likely depends on the strength of household balance sheets. In this paper, we show that excess savings are a useful way to capture this strength or weakness at the business cycle frequency. In the context of the euro area, we find that monetary policy is weaker during periods of higher excess savings. Our finding holds for real and nominal outcomes alike. We rationalize our results with a New Keynesian model in which households value savings to better insure against consumption risk. Through the lens of this model, a high-savings economy is less sensitive to monetary policy shocks than an economy with lower savings, as better insurance leads to a smaller rise in individual consumption risk in contractions. Our findings imply that central banks should track excess savings

¹⁶We recalibrate the steady state in this case by assuming that the central bank adjusts its nominal target to be consistent with the induced rise in the real rate. While the consumption gap between unemployed and employed workers falls, in line with our discussion under intuition, the induced rise in the real rate leads to a small fall in steady-state employment and job-finding rates, but not enough to overturn the improved insurance from higher savings.

¹⁷To keep the model and discussion simple, we did not include the type of frictions in the model that would generate the more hump-shaped and persistent dynamics typically found in the literature. We do not expect them to interfere with our main conclusions. If anything, the higher persistence and gradual buildup could amplify our channel because it works through expectations about economic conditions.

and household balance sheets more generally at a high-enough frequency to gauge the strength of the monetary transmission channel and fine-tune policy decisions.

Appendix A. Additional Data Description

In this section, we elaborate on the savings data used for our empirical analysis and our method for measuring stocks of excess savings.

A.1 *Savings Data*

We collect quarterly household consumption and savings data for each economy in our analysis from national accounts data. Our definition of savings is gross household savings, which are the sum of net household savings and consumption of fixed capital. We define gross nominal household disposable income as the sum of gross household savings and final household consumption. The gross household savings rate is then defined as the following:

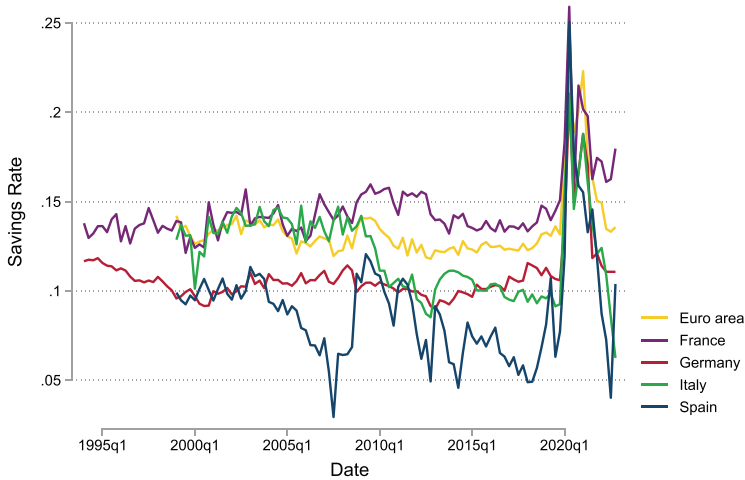
$$\text{Savings rate} = \frac{\text{Gross household savings}}{\text{Gross household disposable income}}.$$

We follow this approach for all countries except Germany. Because of the lack of data on household consumption of fixed capital, we use a net savings concept for Germany where we define its savings rate as the share of net household savings in net household disposable income. Figure A.1 plots the raw savings rates across the economies in our analysis.

A.2 *Measuring Excess Savings*

Following de Soyres, Moore, and Ortiz (2023), we extract a time-varying trend from the savings rate for each country by utilizing the time-series filter proposed in Hamilton (2018). More specifically, for each country we run the following regression of the savings rate on its lags:

$$\begin{aligned} \text{Savings rate}_{t+8} = & \beta_0 + \beta_1 \text{Savings rate}_t + \beta_2 \text{Savings rate}_{t-1} \\ & + \beta_3 \text{Savings rate}_{t-2} + \beta_4 \text{Savings rate}_{t-3} + u_{t+8}. \end{aligned}$$

Figure A.1. Country-Level Savings Rates

Note: Figure A.1 plots the time series of country-level aggregate savings rates used to estimate stock of excess savings.

The residual, u_{t+8} , is our estimate of the deviation of the savings rate from its trend. As stated in the text, we scale this residual by disposable income to obtain a measure, in euros, of the flow of excess savings. We then sum these flows over time to obtain a time-varying measure of the stock of excess savings. We normalize this stock by nominal GDP at each point in time.

Note that for our monthly regressions we specify the one-quarter lag of excess savings as the conditioning variable.

Appendix B. Robustness Results

B.1 ECB Monetary Policy Shock Predictability

To determine whether our monetary policy shocks, which apply the Bu, Rogers, and Wu (2021) approach to the euro area, are predictable using information about the state of the economy, we run the following regression,

$$\text{shock}_t = \alpha + \mathbf{X}'_t \beta + \varepsilon_t, \quad (\text{B.1})$$

Table B.1. ECB Monetary Policy Shock Predictability Regressions

	Monetary Policy Shock (1)	Monetary Policy Shock (2)
Citi Surprise Index	0.000549 (0.00357)	0.000543 (0.00358)
Change in Composite PMI	-0.0166 (0.0727)	-0.0184 (0.0738)
Change in Consumer Sentiment	0.000350 (0.000464)	0.000349 (0.000465)
Two-Quarter-Ahead GDP Growth (SPF)	0.000140 (0.00148)	0.000156 (0.00153)
Scotti Index	-0.00282 (0.00622)	-0.00276 (0.00625)
COVID Dummy		0.00258 (0.00520)
Observations	174	174
R-squared	0.0128	0.0130
<p>Note: Newey-West standard errors are reported in parentheses. + denotes 10 percent significance, ** denotes 5 percent significance, and *** denotes 1 percent significance.</p>		

where \mathbf{X} is a matrix of news variables which includes the Citi economic activity surprise index, the one-month change in the composite PMI (purchasing managers index) for the euro area, the one-month change in consumer sentiment, the two-quarter-ahead forecast of GDP growth from the European Survey of Professional Forecasters (SPF), and the Scotti index of business activity (Scotti 2016). We specify one lag of these variables to ensure that the regressors reflect information available at the time of each ECB meeting and not a reaction to the results of the meeting.

Table B.1 reports the estimated coefficients of regression (B.1). Based on column 1 we do not find any evidence that the ECB monetary policy shocks are predictable on the basis of economic news and expectations. Column 2 reestimates regression (B.1) with an additional control that accounts for COVID. The COVID dummy is set equal to 1 from March 2020 through December 2021. Explicitly

accounting for COVID by specifying the COVID dummy does not change our conclusions.

B.2 Excess Savings Dampen Monetary Policy Effects on Other Measures of Economic Activity

In this section, we document that excess savings dampen the effects of monetary policy on two additional measures of economic activity: real aggregate consumption and industrial production. To show this, we estimate local projections (3) at the quarterly frequency for consumption and the monthly frequency for industrial production, using the percent change in the level of the dependent variable:

$$Y_{i,t+h|t-1} = 100 \times \frac{Y_{i,t+h} - Y_{i,t-1}}{Y_{i,t-1}},$$

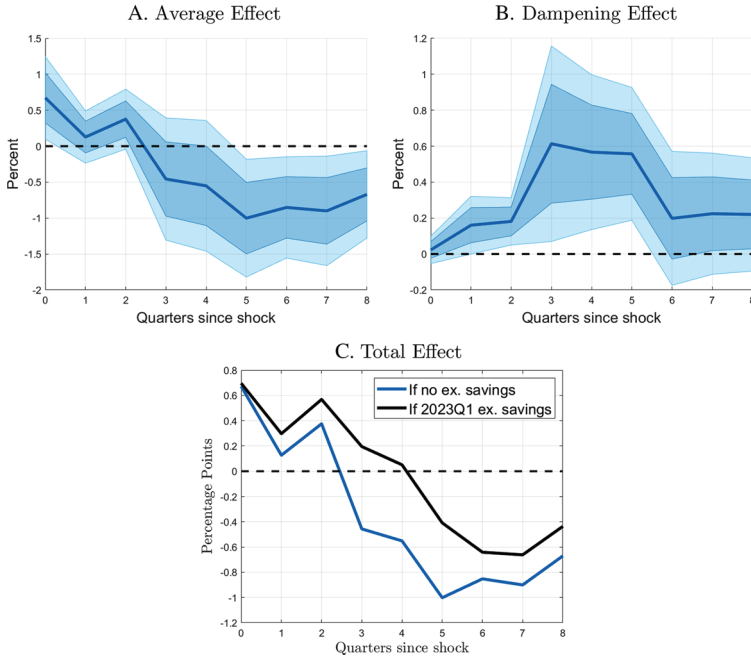
with $Y_{i,t}$ representing consumption and industrial production for economy i at quarter t .

The results are reported in Figures B.1 and B.2. Based on panel A of Figure B.1, we find that a contractionary monetary policy shock reduces real consumption by 1 percent. The decline in consumption, however, is muted when the stock of excess savings is 1 percentage point above the historical average, as shown in panel B of Figure B.1. While, in the absence of excess savings, real consumption falls by nearly 1 percent, panel C of Figure B.1 shows that when excess savings are set equal to their 2023:Q1 level, consumption only declines by about 0.6 percent. Panels A–C of Figure B.2 show qualitatively similar results for industrial production.

B.3 Controlling for COVID Effects

We check whether time-specific changes around COVID drive our results by adding two additional controls to our regression specification: (i) a COVID dummy variable covering March 2020 to December 2021 and (ii) an interaction between the COVID dummy variable and the monetary policy shock. As can be seen from Figures B.3 and B.4, our results do not change much when we specify the COVID dummy and its interaction with the monetary policy shock. This could be partly due to the fact that we already control for some

Figure B.1. Effect of a Contractionary Monetary Policy Shock on Euro-Area Consumption



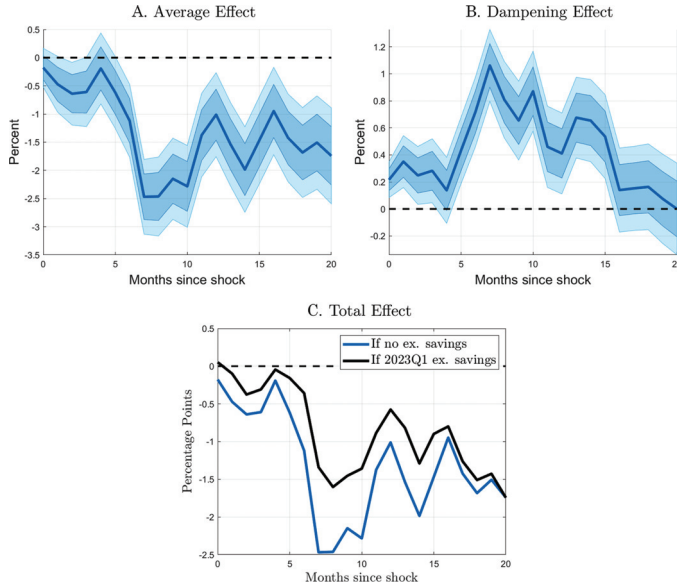
Note: Figure B.1 plots the response of consumption to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect, β_1^h , from local projections (3), while panel B plots the effect conditional on the level of excess savings, β_2^h . Panel C plots the total effect under two different scenarios: (i) when excess savings are equal to 0 (i.e., β_1^h) and (ii) when excess savings are set equal to their 2023:Q1 level based on an average of the countries in our sample. The shaded area reflects the 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay (1998).

state-dependent effects with GDP growth and its interaction with the monetary policy shock.

B.4 Controlling for Bank Balance Sheet Strength

To explore the robustness of our baseline results to bank balance sheet strength, we reestimate our local projections three times using different proxies for bank balance sheet strength: (i) loan-to-deposit ratios, (ii) a measure of the cyclical component of credit to GDP,

Figure B.2. Effect of Tightening Monetary Policy Shock on Euro-Area Industrial Production



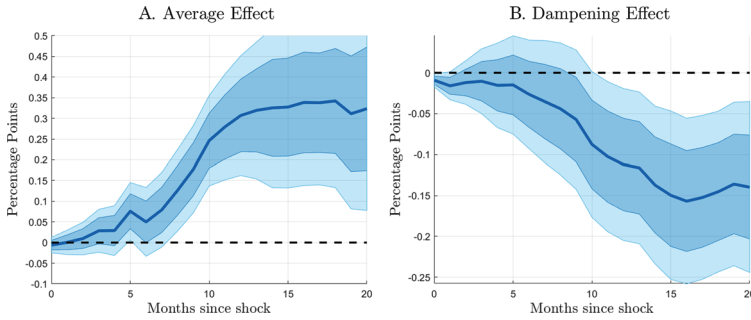
Note: Figure B.2 plots the response of industrial production to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect, β_1^h , from local projections (3), while panel B plots the effect conditional on the level of excess savings, β_2^h . Panel C plots the total effect under two different scenarios: (i) when excess savings are equal to 0 (i.e., β_1^h) and (ii) when excess savings are set equal to their 2023:Q1 level based on an average of the countries in our sample. The shaded area reflects the 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay (1998).

the credit-to-GDP gap, and (iii) bank capital to total assets. A high loan-to-deposit ratio can reflect liquidity risk. Furthermore, the credit-to-GDP gap is regarded as an important variable for banking supervision.¹⁸ Finally, we use bank capital to total assets as another measure of balance sheet strength.¹⁹

¹⁸This variable is frequently used in banking supervision to determine the state of the credit cycle—see, for example, Shin (2013), Drehmann and Tsatsaronis (2014), and Bassett et al. (2015). We obtain it by HP-filtering the credit-to-GDP ratio with a smoothing parameter of 400,000.

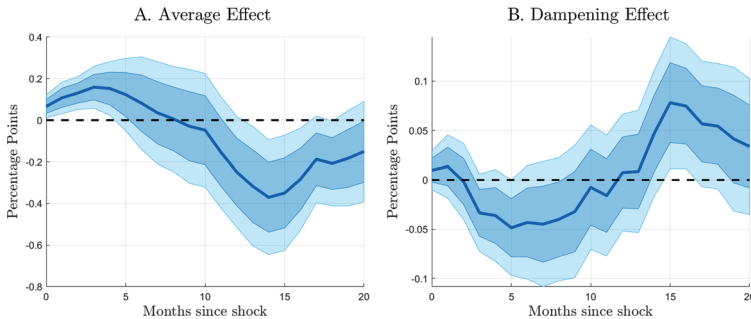
¹⁹Ideally, we would have used tier 1 capital ratios, a common measure of balance sheet strength. However, this variable is unfortunately only available from

Figure B.3. Effect of Contractionary Monetary Policy Shock on Unemployment Rate, Controlling for COVID



Note: Figure B.3 depicts the response of the unemployment rate to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect while panel B plots the effect conditional on the level of excess savings. The shaded area reflects the 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay (1998).

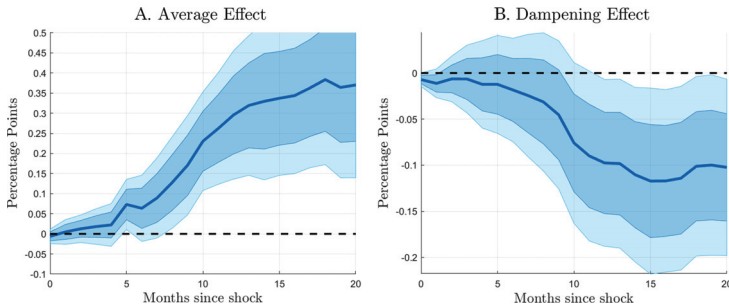
Figure B.4. Effect of Tightening Monetary Policy Shock on Inflation, Controlling for COVID



Note: Figure B.4 depicts the response of the unemployment rate to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect while panel B plots the effect conditional on the level of excess savings. The shaded area reflects the 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay (1998).

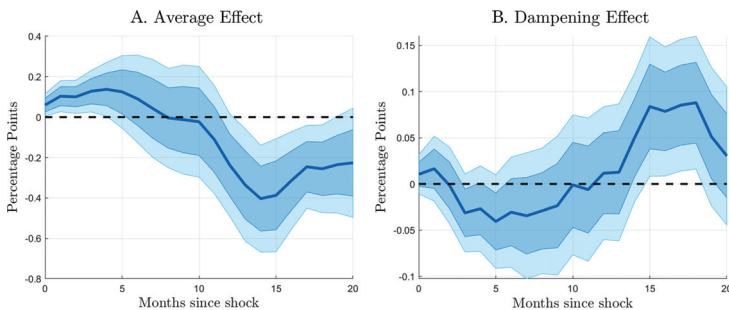
2014 onward at the required frequency. Therefore, using it as a control considerably shortens the time dimension of our panel, reducing statistical power. The three proxies that we use, on the other hand, have longer histories.

Figure B.5. Effect of Contractionary Monetary Policy Shock on Unemployment Rate, Controlling for Loan-to-Deposit Ratios



Note: Figure B.5 depicts the response of the unemployment rate to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect while panel B plots the effect conditional on the level of excess savings. The shaded area reflects 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay (1998).

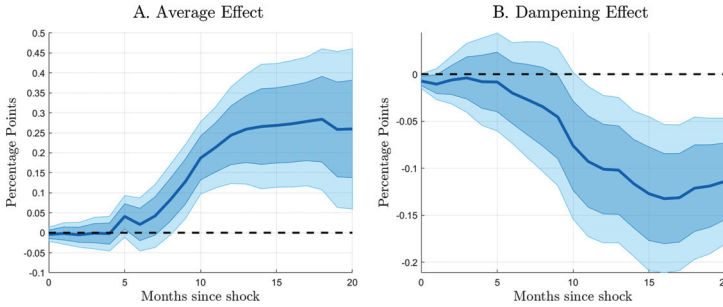
Figure B.6. Effect of Contractionary Monetary Policy Shock on Inflation, Controlling for Loan-to-Deposit Ratios



Note: Figure B.6 depicts the response of the unemployment rate to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect while panel B plots the effect conditional on the level of excess savings. The shaded area reflects the 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay.

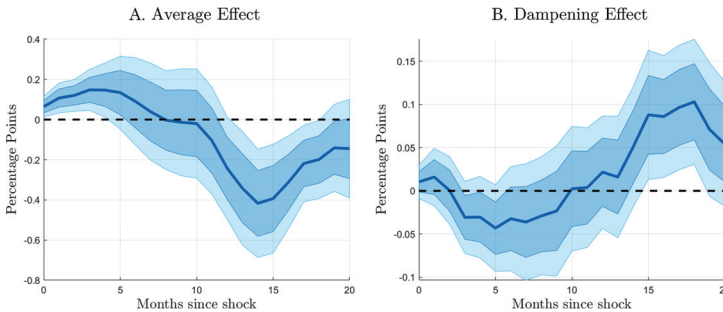
In each of the new regressions, we include lags of the respective balance sheet variable as well as its interaction with the monetary policy shock. Figures B.5 and B.6 plot the results that control for loan-to-deposit ratios. Figures B.7 and B.8 plot the results that

Figure B.7. Effect of Contractionary Monetary Policy Shock on Unemployment Rate, Controlling for Credit-to-GDP Gap



Note: Figure B.7 depicts the response of the unemployment rate to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect while panel B plots the effect conditional on the level of excess savings. The shaded area reflects the 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay (1998).

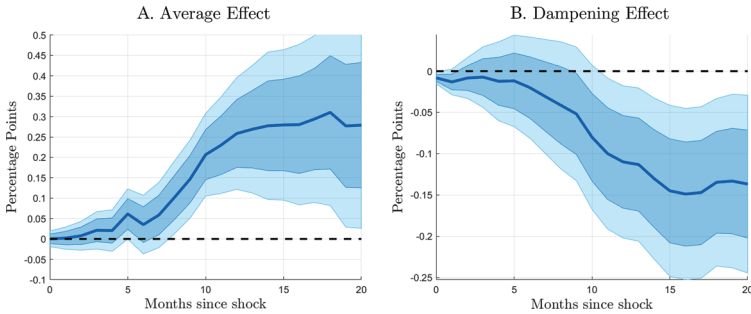
Figure B.8. Effect of Contractionary Monetary Policy Shock on Inflation, Controlling for Credit-to-GDP Gap



Note: Figure B.8 depicts the response of the unemployment rate to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect while panel B plots the effect conditional on the level of excess savings. The shaded area reflects the 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay (1998).

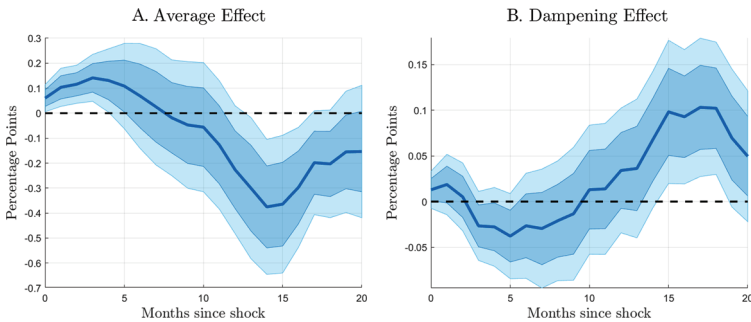
control for the credit-to-GDP gap. Figures B.9 and B.10 plot the results that control for bank capital to total assets. Overall, we find that our results are robust to the inclusion of these different bank balance sheet controls.

Figure B.9. Effect of Contractionary Monetary Policy Shock on Unemployment Rate, Controlling for Bank Capital to Assets



Note: Figure B.9 depicts the response of the unemployment rate to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect while panel B plots the effect conditional on the level of excess savings. The shaded area reflects the 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay (1998).

Figure B.10. Effect of Contractionary Monetary Policy Shock on Inflation, Controlling for Bank Capital to Assets



Note: Figure B.10 depicts the response of the unemployment rate to a monetary policy shock normalized to increase two-year rates by 50 basis points. Panel A plots the unconditional effect while panel B plots the effect conditional on the level of excess savings. The shaded area reflects the 68 percent and 90 percent confidence intervals using standard errors from Driscoll and Kraay (1998).

References

- Abdelrahman, H., and L. E. Oliveira. 2023. "The Rise and Fall of Pandemic Excess Savings." FRBSF Economic Letter No. 2023-11, Federal Reserve Bank of San Francisco.
- Acharya, S., and K. Dogra. 2020. "Understanding HANK: Insights from a PRANK." *Econometrica* 88 (3): 1113–58.
- Aggarwal, R., A. Auclert, M. Rognlie, and L. Straub. 2023. "Excess Savings and Twin Deficits: The Transmission of Fiscal Stimulus in Open Economies." *NBER Macroeconomics Annual 2022*, Vol. 37, ed. M. Eichenbaum, E. Hurst, and V. Ramey, 325–412. University of Chicago Press.
- Aladangady, A., D. Cho, L. Feiveson, and E. Pinto. 2022. "Excess Savings during the COVID-19 Pandemic." FEDS Notes, Board of Governors of the Federal Reserve System, October 21.
- Alpanda, S., E. Granziera, and S. Zubairy. 2021. "State Dependence of Monetary Policy across Business, Credit, and Interest Rate Cycles." *European Economic Review* 140 (November): Article 103936.
- Auclert, A., M. Rognlie, and L. Straub. 2023. "The Trickle Up of Excess Savings." *AEA Papers and Proceedings* 113 (May): 70–75.
- Badinger, H., and S. Schiman. 2023. "Measuring Monetary Policy in the Euro Area Using SVARs with Residual Restrictions." *American Economic Journal: Macroeconomics* 15 (2): 279–305.
- Bassett, W. F., A. Daigle, R. M. Edge, and G. Kara. 2015. "Credit-to-GDP Trends and Gaps by Lender- and Credit-type." FEDS Notes, Board of Governors of the Federal Reserve System, December 3.
- Bauer, M. D., and E. T. Swanson. 2020. "The Fed's Response to Economic News Explains the 'Fed Information Effect'." Working Paper No. 27013, National Bureau of Economic Research.
- Bilbiie, F. O. 2018. "Monetary Policy and Heterogeneity: An Analytical Framework." CEPR Discussion Paper No. 12601.
- Bu, C., J. Rogers, and W. Wu. 2021. "A Unified Measure of Fed Monetary Policy Shocks." *Journal of Monetary Economics* 118 (March): 331–49.
- Challe, E., J. Matheron, X. Ragot, and J. F. Rubio-Ramirez. 2017. "Precautionary Saving and Aggregate Demand." *Quantitative Economics* 8 (2): 435–78.

- Cho, D. 2023. "Unemployment Risk, MPC Heterogeneity, and Business Cycles." *Quantitative Economics* 14 (2): 717–51.
- Cloyne, J., C. Ferreira, and P. Surico. 2020. "Monetary Policy when Households Have Debt: New Evidence on the Transmission Mechanism." *Review of Economic Studies* 87 (1): 102–29.
- de Soyres, F., D. Moore, and J. Ortiz. 2023. "Accumulated Savings During the Pandemic: An International Comparison with Historical Perspective." FEDS Notes, Board of Governors of the Federal Reserve System, June 23.
- Den Haan, W. J., P. Rendahl, and M. Riegler. 2018. "Unemployment (Fears) and Deflationary Spirals." *Journal of the European Economic Association* 16 (5): 1281–1349.
- Drehmann, M., and K. Tsatsaronis. 2014. "The Credit-to-GDP Gap and Countercyclical Capital Buffers: Questions and Answers." *BIS Quarterly Review* (March): 55–74.
- Driscoll, J. C., and A. C. Kraay. 1998. "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data." *Review of Economics and Statistics* 80 (4): 549–60.
- Gornemann, N. M., K. Kuester, and M. Nakajima. 2016. "Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy." Technical Report, Board of Governors of the Federal Reserve System.
- Hamilton, J. 2018. "Why You Should Never Use the Hodrick-Prescott Filter." *Review of Economics and Statistics* 100 (5): 831–43.
- Harding, M., and M. Klein. 2022. "Monetary Policy and Household Net Worth." *Review of Economic Dynamics* 44 (April): 125–51.
- Heathcote, J., and F. Perri. 2018. "Wealth and Volatility." *Review of Economic Studies* 85 (4): 2173–2213.
- Jarociński, M., and P. Karadi. 2020. "Deconstructing Monetary Policy Surprises—The Role of Information Shocks." *American Economic Journal: Macroeconomics* 12 (2): 1–43.
- Jordà, Ò. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–82.
- Miranda-Agrippino, S. 2016. "Unsurprising Shocks: Information, Premia, and the Monetary Transmission."

- Ramey, V. A. 2016. "Macroeconomic Shocks and Their Propagation." *Handbook of Macroeconomics*, Vol. 2, ed. J. B. Taylor and H. Uhlig, 71–162. North-Holland.
- Ravn, M. O., and V. Sterk. 2017. "Job Uncertainty and Deep Recessions." *Journal of Monetary Economics* 90 (October): 125–41.
- Scotti, C. 2016. "Surprise and Uncertainty Indexes: Real-Time Aggregation of Real-Activity Macro-surprises." *Journal of Monetary Economics* 82 (September): 1–19.
- Shin, H. S. 2013. "Procyclicality and the Search for Early Warning Indicators." Working Paper No. 13/258, International Monetary Fund.
- Tenreyro, S., and G. Thwaites. 2016. "Pushing on a String: US Monetary Policy is Less Powerful in Recessions." *American Economic Journal: Macroeconomics* 8 (4): 43–74.