

Monetary Policy and Liquidity Constraints: Evidence from the Euro Area[†]

By MATTIAS ALMGREN, JOSÉ-ELÍAS GALLEGOS,
JOHN KRAMER, AND RICARDO LIMA*

We quantify the relationship between the response of output to monetary policy shocks and the share of liquidity-constrained households. We do so in the context of the euro area, using a Local Projections Instrumental Variables estimation. We construct an instrument for changes in interest rates from changes in overnight indexed swap rates in a narrow time window around ECB announcements. Monetary policy shocks have heterogeneous effects on output across countries. Using micro data, we show that the elasticity of output to monetary policy shocks is larger in countries that have a larger fraction of households that are liquidity constrained. (JEL E23, E43, E52, F33, F45, G51)

In 2016, 30 percent of households in Germany reported that they could not meet an unexpected, immediate expense of 985 euros. At the same time, 40 percent of Italian households reported that they would be unable to meet an unexpected expense of 800 euros.¹ Figures like these suggest that a significant portion of households hold few liquid assets, which potentially makes them vulnerable to unexpected shocks to the economy. These households have received special attention recently, especially in monetary economics.

*Almgren: Institute for International Economic Studies (IIES), Stockholm University (email: mattias.almgren@iies.su.se); Gallegos: Banco de España (email: jose.elias.gallegos@bde.es); Kramer: Department of Economics, University of Copenhagen (email: john.kramer@econ.ku.dk); Lima: IIES, Stockholm University (email: ricardo.lima@iies.su.se). Giorgio Primiceri was coeditor for this article. This paper was previously circulated under the title “Monetary Policy and the Wealthy Hand-to-Mouth.” We are grateful to Tobias Broer, Per Krusell, Kurt Mitman, and Kathrin Schlafman for their advice and support. Further, we would like to thank four anonymous referees, Adrien Auclert, Martín Beraja, Mitch Downey, Xavier Gabaix, Alessandro Galesi, Pierre-Olivier Gourinchas, John Hassler, Greg Kaplan, Peter Karadi, Karin Kinnerud, Kasper Kragh-Sørensen, Jesper Lindé, Benjamin Moll, Haroon Mumtaz, Morten Ravn, Ricardo Reis, Giovanni Ricco, Federica Romei, Maria Sandström, David Schönholtzer, József Sigurdsson, Xueping Sun, Javier Vallés, Has van Vlokhoven, Iván Werning, and seminar participants at the IIES, Queen Mary University of London, Stanford University, Stockholm University, Stockholm School of Economics, Università Bocconi, and Yale University for useful feedback and comments. Björn Hagströmer generously aided us in accessing the swap rate data. José-Elías Gallegos and John Kramer gratefully acknowledge funding from the Tom Hedelius foundation. This paper was included as Chapter 1 in José-Elías Gallegos’s dissertation and as Chapter 2 in John Kramer’s dissertation, both at the IIES, Stockholm University. This paper uses data from the Eurosystem Household Finance and Consumption Survey. The results published and the related observations and analysis may not correspond to the results or analysis of the data producers.

[†]Go to <https://doi.org/10.1257/mac.20200096> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

¹Responses according to the European Union Survey of Income and Living Conditions. The monetary values represent the country-specific at-risk-of-poverty thresholds, defined as 60 percent of the national median equivalized disposable income after social transfers.

While theoretical research has shown that heterogeneous agent models that include constrained agents can have different policy implications than their representative agent counterparts, empirical evidence on how heterogeneity matters for the transmission from monetary policy to output is scant.² In this paper we provide such evidence, showing that a higher share of liquidity-constrained households in a country is associated with a stronger output response to a monetary policy surprise.

We focus on the euro area, where member countries have been exposed to the common monetary policy conducted by the European Central Bank (ECB) since the introduction of their shared currency. However, because of long-standing national idiosyncrasies and slow convergence, they still differ along many dimensions—including the share of liquidity-constrained households, as we show. Since we choose this bird's-eye view, we can conduct standard monetary policy analysis while taking account of wealth and income heterogeneity and its influence on output responses.

First, we estimate output impulse response functions (IRFs) at monthly frequencies for each country to the same monetary policy shocks, relying on the Local Projection (LP) approach pioneered by Jordà (2005).³ Because of endogeneity concerns between policy rate changes and output responses, we augment the LP estimation with an instrumental variable (IV) framework (Stock and Watson 2018).⁴ We use high-frequency movements in Overnight Indexed Swap (OIS) rates in a 45-minute time window around ECB policy announcements as an instrument for monetary policy surprises. Because OIS rates are forward-looking interest rate derivatives, large rate movements during the window imply that the ECB's announcement was not in line with market expectations. The identifying assumption is that this measure is uncorrelated with other shocks to output.

In the second part of the paper, we incorporate the income and asset dimensions by relating the IRFs to the share of liquidity-constrained households in each country. The idea is that a higher fraction of households less able to smooth the income fluctuations caused by monetary policy shocks may lead to a stronger aggregate output response in a country. While it is not possible to measure the fraction directly, we approximate it by classifying households in the Household Finance and Consumption Survey (HFCS) as Hand-to-Mouth (HtM) or non-HtM according to a measure proposed by Kaplan, Violante, and Weidner (2014). They show that such a measure is strongly correlated with estimates of marginal propensities to consume (MPC). Since the HFCS can only provide data on recent years, we complement it with data from the European Union Survey on Income and Living Conditions (EU-SILC), which has been conducted since 2005. Participating households are

²See, e.g., Bilbiie (2008) for an early theoretical contribution in a two-agent setting or Auclert (2019) and Hagedorn et al. (2019) for a setting with fully heterogeneous agents.

³Mandler, Scharnagl, and Volz (2016) investigate a similar question using a Bayesian VAR model for the four largest economies in the euro area: Germany, Italy, Spain, and France. Altavilla, Giannone, and Lenza (2016) investigate heterogeneous effects of Outright Monetary Transactions (OMT) on the same countries, similarly employing a VAR framework.

⁴As a robustness check to our main empirical framework, we construct an instrumented Global VAR model (GVAR) based on Georgiadis (2015) and Burriel and Galesi (2018). We build a more structural—and restricted—setting than the LPIV, more similar to the widespread VAR estimation in the literature, identifying monetary responses in a GVAR setting using exogenous instruments. To our knowledge, we are the first to estimate such an instrumented GVAR. We find similar results.

asked whether they could finance an unexpected financial expense, from which we infer whether they are financially constrained. Both surveys point to large variation across countries in the share of constrained consumers, and the pattern is broadly consistent over time.

Our first finding is that, in line with previous literature, output responses to a common European monetary policy surprise are not homogeneous across countries. There is significant heterogeneity in terms of cumulative impact and peak values. Secondly, all of our measures of the fraction of liquidity-constrained households are significantly correlated with the strength of the IRFs. On average, countries with higher fractions of liquidity-constrained households exhibit stronger cumulative output responses and bigger peak responses to an unexpected interest rate change. For the measure constructed according to Kaplan, Violante, and Weidner (2014), we show that the results are driven by the “wealthy HtM,” i.e., households with low levels of liquid wealth but positive and possibly large levels of illiquid wealth. In addition, we calculate aggregate output IRFs for a constrained and a less constrained group of countries. The two responses are significantly different at most horizons, with the more constrained countries reacting more strongly to the common shock.

The results we present are important for several reasons. First, our findings suggest that heterogeneity in the composition of household balance sheets across countries affects the transmission of monetary policy to their economies. The finding that a higher share of low-liquidity households amplifies the output response to an unexpected interest rate change can guide future theoretical and quantitative work on monetary policy in a Heterogeneous Agent New Keynesian framework. Understanding the reasons for the differences we uncover is crucial in order to calibrate future policies. Second, we show that LP methods can be used to estimate the impact of monetary policy for countries within a currency union. Lastly, our results are robust across different specifications of liquidity constraints, corroborating the measure put forth by Kaplan, Violante, and Weidner (2014).

Our research is related to several strands of literature. There is a large body of research that performs cross-country monetary policy analysis. An early example is Gerlach and Smets (1995), who perform a structural VAR analysis of the G7 countries and find that responses to country-specific monetary policy shocks are similar. Mandler, Scharnagl, and Volz (2016), using a large Bayesian VAR, show that output in Spain is less responsive to monetary policy compared to Germany, France, and Italy, while prices in Germany respond most within this set of countries. Few papers estimate IRFs for multiple countries and try to investigate the transmission mechanism of monetary policy by relating their findings to country characteristics. Two recent examples, both of which use a Global VAR (GVAR) method, are Georgiadis (2015) and Burriel and Galesi (2018). Both papers find heterogeneous responses of real GDP across countries and explain some of the variation with wage rigidities and the fragility of the banking sector. Calza, Monacelli, and Stracca (2013) provide evidence that in countries where the use of flexible mortgage rates is more prevalent, responses to monetary policy shocks are stronger, and Corsetti, Duarte, and Mann (2018) find that the responses of output and private consumption are larger in countries where home ownership rates are higher. We try to account for previous findings by conducting several robustness checks.

To our knowledge we are the first to use OIS rates as an instrument to identify a cross-country LP estimation in the euro area. Kuttner (2001), Nakamura and Steinsson (2015), and Gertler and Karadi (2015) use high-frequency movements in federal funds futures rates in a short window around the Federal Reserve's policy announcements to identify monetary policy surprises in the United States. In the European context, there are no financial instruments equivalent to federal funds futures, which has led researchers to employ high-frequency movements in OIS rates instead. Ampudia and Van den Heuvel (2017) and Jarociński and Karadi (2020) construct monetary policy shocks from movements in these derivatives.

The empirical results in this paper tie in with the results from theoretical Two-Agent New Keynesian (TANK) models such as those in Bilbiie (2008); Galí, Lopez-Salido, and Valles (2007); and Bilbiie (2020b), as well as richer models by Gornemann, Kuester, and Nakajima (2016); Werning (2015); Auclert (2019); and Hagedorn et al. (2019). As laid out by Bilbiie (2020a), one result these models have in common is that whether aggregate shocks have bigger or smaller effects on aggregate consumption compared to the representative agent framework, is ambiguous. In a model that combines the tractability of TANK models with the most important elements of heterogeneous agent models, Bilbiie (2020a) shows that the output response to shocks is amplified if the income elasticity of constrained agents with respect to aggregate income is larger than one. He refers to this case as procyclical income inequality, a channel that is strengthened if a larger fraction of agents is constrained.⁵ This is in line with our empirical findings, which can guide future modeling efforts aimed at understanding the interaction between aggregate and distributional outcomes in response to shocks.

Lastly, our findings imply that it is important to treat liquid and illiquid assets separately when describing the wealth distribution of an economy. This is in support of the view that wealthy households can have high marginal propensities to consume, as pointed out by Kaplan and Violante (2014); Kaplan, Violante, and Weidner (2014); and Kaplan, Moll, and Violante (2018).⁶

The paper proceeds as follows. In Section I we describe our identification strategy and how we estimate country-specific local projections and present the resulting IRFs. Section II discusses how we construct the proxies for the fraction of liquidity-constrained households across countries. Section III relates the IRFs to our measures of the fraction of liquidity-constrained households across countries. Section IV concludes.

⁵In models that focus on the cyclical risk of income risk (e.g., Werning 2015), amplification of aggregate shocks is caused by an increase in the probability of becoming constrained for the unconstrained, which leads the latter to save more and consume less. Our empirical analysis, however, focuses mainly on the level of the HtM shares (as opposed to their changes) and is therefore more closely related to Bilbiie (2020a).

⁶Using data from Norwegian lottery winners, Fagereng, Holm, and Natvik (2021) find that households at the highest liquidity quartile have a significantly lower MPC than households at the lowest liquidity quartile.

I. Effects of Monetary Policy Shocks on Output

A. Identifying Monetary Policy Shocks

In order to estimate the effects of monetary policy on a variable of interest, we need to identify unexpected deviations from an interest rate rule. To identify these in the United States, researchers have used high-frequency movements in federal funds futures in a narrow time window around announcements by the Federal Reserve (Kuttner 2001; Nakamura and Steinsson 2015). More recently, Ampudia and Van den Heuvel (2017) and Jarociński and Karadi (2020) have applied this approach to European data using Overnight Indexed Swap (OIS) rate movements around ECB announcements. These derivatives are traded over the counter between two parties exchanging a fixed interest rate for the floating Eonia overnight interest rate, both on a notional principal, for a prespecified amount of time. Since the principal is not exchanged at any time and the contracts are highly collateralized, there is only minimal counterparty credit risk. When the contract ends, the difference between the fixed interest accrued on the principal and the interest accrued on the principal by investing it at the overnight interest rate every day is calculated, and the contract is cash settled.⁷

We follow the literature and use changes in Eonia OIS during a short time window around the ECB's monetary policy announcement and the subsequent press conference as our instrument, like Jarociński and Karadi (2020).⁸ On days when the governing council of the ECB decides the policy rate for the euro area, the decision is communicated to the public via a press statement at 13:45 CET and motivated during a press conference chaired by the president and vice president at 14:30 CET. We construct a time series encompassing all such monetary policy announcements by the ECB starting in December 1999.⁹ Figure 1 displays the OIS rate on July 5, 2012. The first window starts 15 minutes before the press release and ends 30 minutes after. The second window starts 15 minutes before the beginning of the press conference and ends 30 minutes after. To construct our instrument, on each announcement date we calculate the change in the average OIS rates of the pre- and post windows for both the press statement and the press conference and then sum the two.

The OIS can be viewed as an indicator for expectations about future overnight interest rates in the European interbank market. Hence, a significant change in the OIS rates shortly after an ECB monetary policy announcement implies that the content of the announcement was at least partly unexpected. The identifying assumption is that there is no other information released during the time window that is systematically related to the policy decision and that the market has access to the same information about economic fundamentals as the ECB. As pointed out by Jarociński

⁷For a detailed discussion of similarities between federal funds futures and overnight indexed swaps, see Lloyd (2018).

⁸We obtain time series on OIS rates at the minute frequency from Datastream. For more information, see online Appendix B.B3.

⁹To construct monetary shocks starting from January 2000, we start collecting movements in OIS rates from December 1999, due to the way we construct our instrument (see below).

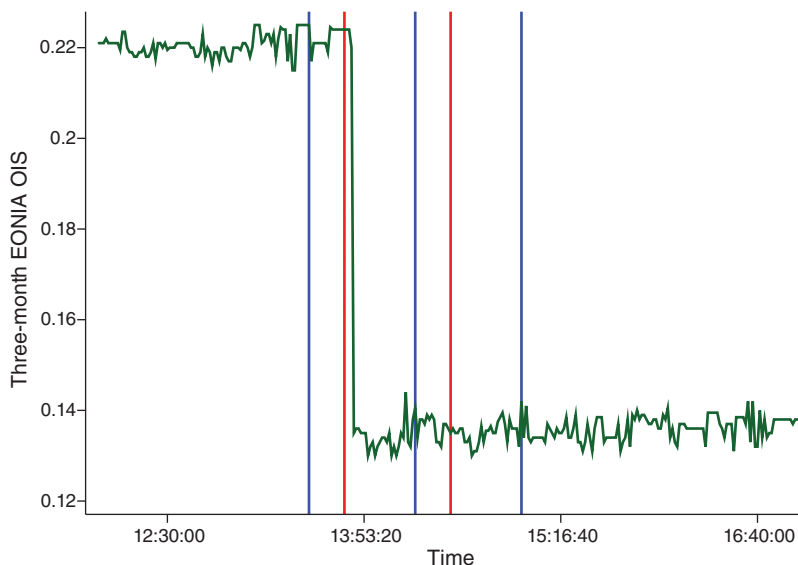


FIGURE 1. OVERNIGHT INDEXED SWAP RATES ON JULY 5, 2012

Notes: This figure shows the time series of the three-month EONIA Overnight Indexed Swap for July 5, 2012. The blue lines represent the borders of our measurement windows. The red lines indicate the policy events; i.e., the ECB's press release at 13:45 CET and the start of the press conference at 14:30 CET. The first pre-window runs from 13:30 to 13:44 CET and the first post window is active between 13:45 and 14:14 CET. Then a second pre-window runs from 14:15 to 14:29 CET and the second post window is active between 14:30–14:59 CET.

and Karadi (2020), many of the Bank of England's announcement dates coincide with announcement dates of the ECB, with policy statements released at 13:00 CET and 13:45 CET, respectively. This makes the high-frequency approach especially important in our setting. The use of instruments measured at the daily frequency would confound the effects of the former and the latter.

We use the three-month OIS rate obtained from Datastream. To convert the instrument series obtained in this way to monthly frequency, we follow Gertler and Karadi (2015). Because the announcement days are at different times during each month, we weigh each observation according to when in a month it occurred. Let a_d be the cumulative shock series at day d in the month, which evolves in the following way:

$$a_d = \begin{cases} a_{d-1} + \Delta f_d & \text{if announcement at day } d \\ a_{d-1} & \text{otherwise,} \end{cases}$$

where Δf_d is the change in the OIS rates calculated as described above. We then weigh the series according to

$$F_t = \frac{1}{D_m} \sum_{d \in m} a_d,$$

where D_m is the number of days in month m . Finally, the instrument for each month t is

$$Z_t = F_t - F_{t-1}.$$

B. The Effect of Monetary Policy Shocks on Output

We follow Jordà (2005) and Stock and Watson (2018) and estimate the response of output to monetary policy shocks using the local projections instrumental variable (LPIV) method, employing the instrument discussed in the previous section. Impulse responses for each country n are constructed from the sequence $\{\beta_n^h\}_{h=0}^H$ from the estimated equations

$$(1) \quad y_{n,t+h} - y_{n,t-1} = \alpha_n^h + \beta_n^h \hat{i}_t + \sum_{j=1}^p \Gamma_{n,j}^h X_{t-j} + u_{n,t+h}, h = 0, \dots, H,$$

where y_n is log of output in country n , X is a set of control variables and \hat{i} are the fitted values from the first-stage regression¹⁰

$$(2) \quad i_t = c + \rho Z_t + \sum_{j=1}^p D_j^h X_{t-j} + e_t.$$

As a benchmark we set the number of lags to $p = 3$ and the horizon of the impulse responses to $H = 36$.¹¹ In all specifications we include the interest rate (i), the instrument Z , aggregate euro-area output, and the price level in our set of control variables, X .^{12,13} Notice that equation (2) resembles a standard Taylor rule for the ECB: the current interest rate depends on lags of euro-area output and inflation, plus lags on the interest rate itself.¹⁴

Our dependent variable, monthly GDP, is measured as the logarithm of real GDP. Given that GDP is only available at quarterly frequency, we follow Chow and Lin (1971) to interpolate real GDP into a monthly frequency.¹⁵ We use the Euro Overnight Index Average (EONIA) as the monetary policy rate and the logarithm of the deseasonalized Harmonized Index of Consumer Prices (HICP) as the measure of the aggregate price level. We use data from January 2000 to December 2012, capturing the initial stages of the adoption of the euro and ending during the year when the interest rate hit the zero lower bound.

Figure 2 presents impulse responses of real GDP for each country in our sample to an expansionary shock of one standard deviation in our instrument, following Jarociński and Karadi (2020). The IRFs are represented by the blue lines. Following

¹⁰For a detailed description of the data series used, we refer the reader to online Appendix B.B6.

¹¹We also estimate specifications in which the number of lags is allowed to vary across the countries using the Akaike Information Criteria (AIC). Doing so leaves the results unaltered and, therefore, for simplicity, we choose the same number of lags for all countries.

¹²As pointed out by Ramey (2016), the construction of the instrument as in Gertler and Karadi (2015) introduces autocorrelation into the instrument, invalidating our identifying assumptions. To alleviate this problem, we include the instrument in addition to the other control variables in X .

¹³Removing the set of lagged control variables in (2), especially the interest rate, leads to very low F -statistics. Since the interest rate is persistent, contemporaneous shocks account for only a small part of its variance. Furthermore, of the contemporary shocks, the monetary policy shock is only a fraction. Therefore, the explanatory power of the instrument alone on the interest rate can be expected to be fairly low Stock and Watson (2018). The first state F -statistic in our benchmark specification is 17.42.

¹⁴As a robustness check, we conduct the same exercise including country-specific lags in equation (2), leading to country-specific first-stage regressions and country specific \hat{i} s. The results are reported in Figure A2 in the Appendix.

¹⁵For each euro-area country as well as the aggregate euro area, we use monthly data for industrial production, retail trade, and unemployment to construct monthly series for real GDP.

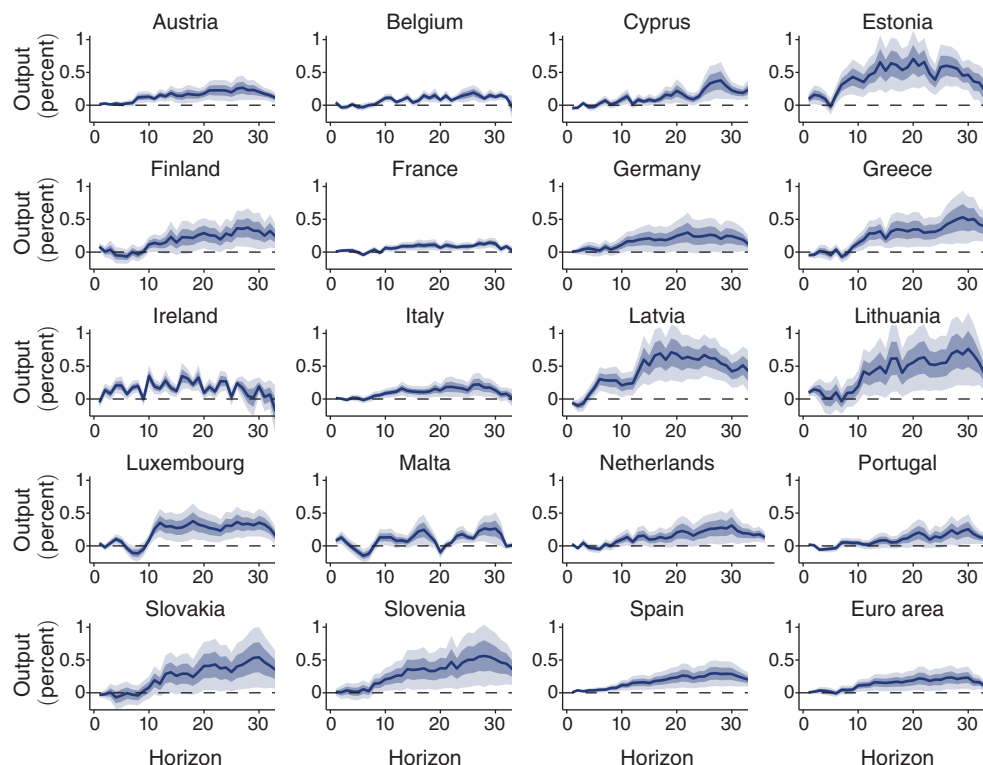


FIGURE 2. IMPULSE RESPONSES FOR OUTPUT IN EURO AREA COUNTRIES—LPIV

Notes: This figure shows impulse responses of real GDP to an expansionary monetary policy shock of one standard deviation. For each euro-area country, the response is estimated using LPIV (equation (1)). The solid blue lines represent the IRFs produced by our preferred specification (see text for details). The dark blue and light blue shaded areas represent one- and two-standard-deviation confidence bands, respectively, constructed using Newey-West estimators.

Stock and Watson (2018), we construct one- and two-standard-deviation confidence bands that surround the point estimates, using Newey-West estimators.¹⁶

The estimated impulse response functions reveal that expansionary monetary policy shocks cause output to increase in most countries. Output increases significantly after a little less than a year, with the maximum impact most often occurring later. The response of aggregate euro-area output, for example, reaches its peak of 0.23 percentage points after 27 months.

There is considerable heterogeneity in the magnitudes of the responses, both in peak and cumulative effects. Moreover, the initial impacts of monetary policy shocks seem to be, on average, small and often not statistically different from zero.

Using the results from Georgiadis (2015) as a proxy for the VAR counterpart of our analysis, we find that the peak values are strongly correlated for the subset of

¹⁶The local projection impulse responses for prices are presented in Figure A3 in the Appendix.

countries that overlap with his sample, with a correlation coefficient of approximately 0.84.¹⁷ Given that the relative positions of countries is important for the analysis in the upcoming section, we consider it to be reassuring that our estimates are in line with the findings in Georgiadis (2015).

We proceed now by relating the strength of the output responses to the share of liquidity-constrained households in each country.

II. Measuring Financial Constraints

Bilbiie (2020b) describes a TANK economy in which household heterogeneity is collapsed to being either financially constrained or not. Taking this idea to the data, our aim is to construct variables that measure the degree of financial constraints in a given country. To do so, we rely on the Eurosystem Household Finance and Consumption Survey (HFCS) and the European Union Statistics on Income and Living Conditions (EU-SILC). In the next subsections, we describe these datasets and the construction of our measures for financial constraints used in the subsequent analysis.

A. The Household Finance and Consumption Survey

The HFCS is conducted by the Household Finance and Consumption Network (HFCN), tasked by the Governing Council of the ECB. The survey is modeled after the US Survey of Consumer Finances and is harmonized across the euro area, set up to collect micro data on household finances (Honkkila and Kavonius 2013). It contains data from interviews with over 84,000 households. Three waves have been conducted, with data releases in 2013, 2016, and 2020. In our main analysis, we rely on data from the second wave.

In approximating the share of households that have high MPC, we follow Kaplan, Violante, and Weidner (2014).¹⁸ A household is categorized as HtM if its liquid wealth is smaller than a certain share of monthly income. In their set of countries, the share of HtM households is between 20 to 35 percent.¹⁹

Let m_i denote liquid assets, a_i denote illiquid assets, y_i denote income, and \underline{m}_i be a credit limit for household i .²⁰ We categorize a household as HtM if

$$(3) \quad 0 \leq m_i \leq \frac{y_i}{2}$$

or if

$$(4) \quad m_i \leq 0, \text{ and } m_i \leq \frac{y_i}{2} - \underline{m}_i.$$

¹⁷ Georgiadis (2015) estimates responses for Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, the Netherlands, Portugal, Slovakia, Slovenia, and Spain.

¹⁸ Kaplan, Violante, and Weidner (2014) find that households categorized as HtM according to their measure have an estimated MPC of more than twice that of non-HtM households.

¹⁹ This includes the United States, Canada, Australia, the United Kingdom, Germany, France, Italy, and Spain.

²⁰ See online Appendix B.B4 for details about the classification of assets as liquid or illiquid.

The credit limit \underline{m}_i is set to be the household's monthly income, capturing the possibility of spending using a credit card and repaying the debt once a month. For our sample, the fraction of households who are categorized according to equation (4) is small.

We further divide households into wealthy and poor HtM, following Kaplan, Violante, and Weidner (2014). A household is categorized as wealthy HtM if, in addition to fulfilling one of the conditions in equations (3) and (4), it has positive illiquid wealth:

$$(5) \quad a_i > 0.$$

If a household satisfies either one of equations (3) or (4), but not the condition in equation (5), we label that household as poor HtM.²¹

Figure 3, panel A plots the total fraction of HtM households as well as the split between wealthy and poor. The cross-country variation is striking, with the fraction of HtM households ranging from just above 10 percent in Malta to almost 65 percent in Latvia. In most countries (exceptions being Austria, France, Germany, and Ireland) the fraction of wealthy HtM households exceeds the fraction of poor HtM households, which is in line with the findings in Kaplan, Violante, and Weidner (2014).²²

A concern with the measure described above is that households are interviewed at different points during the month or year. If there are systematic differences across countries in when households are interviewed, this could lead to biased estimates. To combat this, we construct a second proxy for a household's MPC that relies on the past year's income and expenses.

In the HFCS questionnaire, households are asked if, over the last 12 months, their expenses (i) exceeded income, (ii) were about the same as income or (iii) were less than income. A household in categories (i) or (ii) is likely more sensitive to unexpected shocks than one in category (iii) and is therefore likely to have a higher MPC. We compute the fraction of households whose expenses were about the same as or exceeded income and label households that fulfill this criteria as Potentially Financially Vulnerable type 1 (PFV1).

The fractions are presented in Figure 3, panel B. Again, there is heterogeneity across countries and the average, indicated by the vertical line, is above 60 percent. For all countries, the share of PFV1 households is higher than the HtM share. We consider this statistic an upper bound for the fraction of households who have high MPC, as it disregards the possibility that they might have substantial amounts of liquid assets. The correlation between PFV1 and HtM is 0.68, which we see as encouraging.

²¹ For a discussion about the theory behind this classification scheme, we refer the reader to Kaplan, Violante, and Weidner (2014).

²² Most households that are classified as W-HtM own the residence in which they live. The data show that in most countries the majority of households that are classified as W-HtM do not have a mortgage; the fraction varies between 0.12 and 0.67 with a mean (median) of 0.34 (0.34). This fraction appears to be negatively correlated with the fractions of W-HtM across countries. See online Appendix B.B5 for more details.

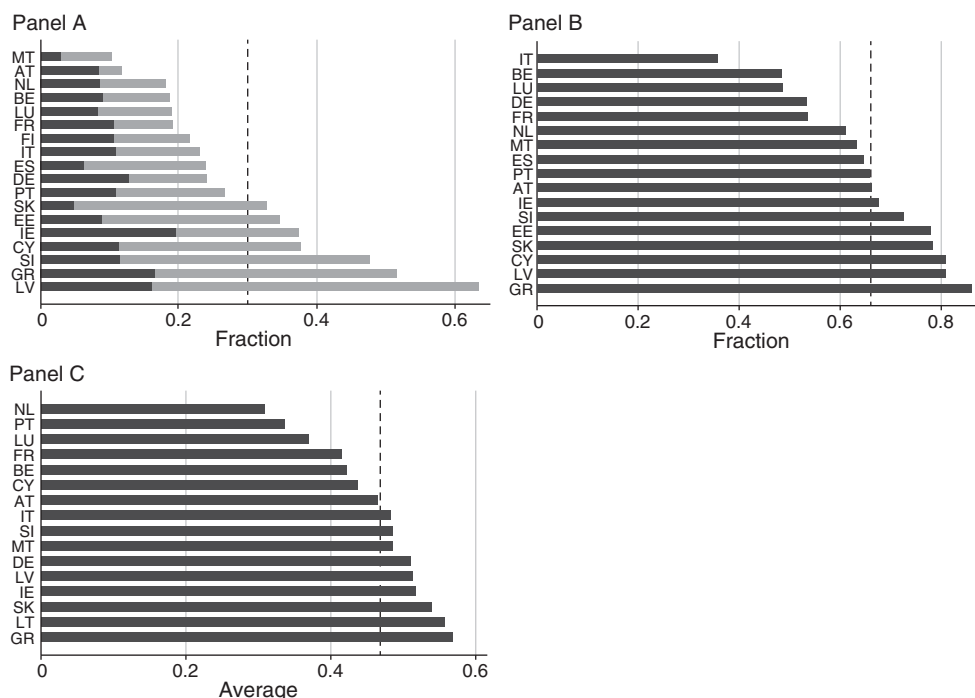


FIGURE 3. HFCS PROXIES FOR HAND-TO-MOUTH STATUS

Notes: Panel A shows the fraction of households that are classified as HtM in each country. The total fraction, given by the total length of each bar, is divided up into two parts: poor (black) and wealthy (gray). The vertical line indicates the unweighted average of total HtM in our sample of countries. We do not have data for Lithuania. Panel B shows the fraction of households that have had expenses over the last 12 months that were “about the same as” or exceeded their income over the same period. The total fraction is given by the total length of each bar. The vertical line indicates the average of the fractions in our sample of countries. Data are missing for Finland and Lithuania; hence, they are not included in the figure. Panel C shows the average of fractions of lottery winnings in each respective country that households would spend over the next 12 months. See text for a more detailed description. Data do not exist for Estonia, Finland, and Spain. For panels A and B, we use data from the second wave of the Eurosystem Household Finance and Consumption Survey (HFCS). For panel C we use data from the third wave of the HFCS.

The most recent wave of the HFCS introduces a new question that attempts to capture MPC in a more direct way. Households are asked what percentage of a hypothetical lottery win they would spend over the next 12 months.²³ Within each country, we compute the average of these reported MPC across all households. Figure 3, panel C presents the resulting averages, and we can see that there is considerable variation across the countries. The correlation between this measure and our HtM measure is 0.49.

²³The question reads, “Imagine you unexpectedly receive money from a lottery, equal to the amount of income your household receives in a month. What percent would you spend over the next 12 months on goods and services, as opposed to any amount you would save for later or use to repay loans?”

B. The European Union Statistics on Income and Living Conditions Survey

The sample period for the two measures derived above coincides with the end of the European sovereign debt crisis. To ensure that this is not driving our results, we construct two additional variables from the European Union Statistics on Income and Living Conditions (EU-SILC) questionnaire. The EU-SILC is a yearly survey with the objective of measuring income, poverty, social exclusion, and living conditions in the European Union and is executed by the national statistical authorities. At its introduction in 2003 it covered seven countries, and since 2005 it has covered all the countries in our sample, with a sample size of close to 90,000 households.²⁴ Because of its early inception, the survey allows us to construct proxies for the share of households with high MPC with data from before the Great Recession.

First, we use a question about whether a household, out of its own resources, would be able to cover a hypothetical, unexpected, required financial expense equal to the national monthly at-risk-of-poverty threshold.²⁵ Households who expect not to be able to do so are likely to have high MPC out of transitory income shocks. We take the share of households unable to face an unexpected expense as a percentage of all households in 2005 and label it Potentially Financially Vulnerable type 2 (PFV2). Figure 4, panel A displays the variable across countries. Although it is calculated using a different survey and a different sampling period, the correlation coefficient between PFV2 and HtM is 0.67.

We construct one more variable using the EU-SILC survey from 2005. In the survey, households are asked if they were unable to pay utility bills during the last year on time (have been in arrears) due to financial difficulties.²⁶ We assume that households to whom this applies will consume a large share of an unexpected income shock and therefore classify these households as having high MPC and all others as having low MPC. The share of the former in the population is Potentially Financially Vulnerable type 3 (PFV3). The cross-sectional distribution of PFV3 is shown in Figure 4, panel B. For all countries, this measure is the lowest. Intuitively, all other indicators measure the potential of not being able to make ends meet for a household, while PFV3 is the share of households who are already behind on making payments. Therefore, it can be viewed as the strictest proxy among the ones presented in this section, and we view it as the lower bound of households with high MPC. The correlation between PFV3 and the HtM measure is 0.80.

III. Liquidity-Constrained Households and Monetary Policy Effectiveness

A. Results

The results in Section IB indicate that the countries in our sample do not respond homogeneously to monetary policy shocks. We proceed to link this finding with

²⁴The sample size for the whole survey is about 130,000 households. The figure in the text refers to the 19 countries in our sample.

²⁵The at-risk-of-poverty threshold is defined as 60 percent of the national median equivalized disposable income after social transfers.

²⁶Utility bills include heating, electricity, gas, water, etc.

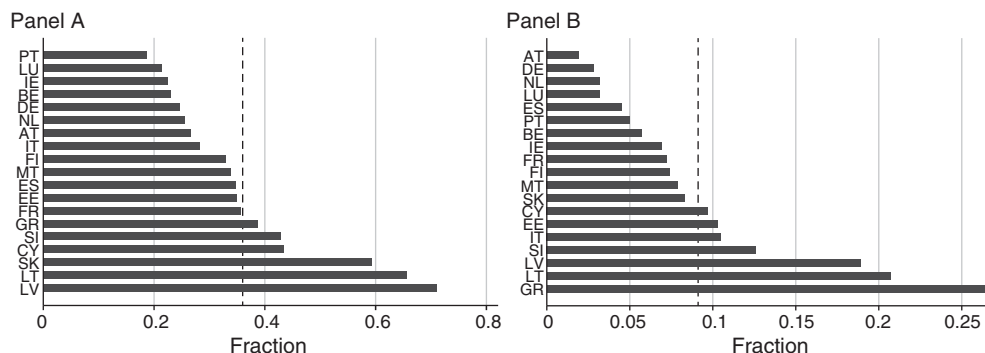


FIGURE 4. PFV2 AND PFV3

Notes: Panel A shows the fraction of households that believe that they are unable to face unexpected expenses with the use of their own resources (PFV2). The fraction is given by the length of each bar. The vertical line indicates the average of the fractions in our sample of countries. Data are from European Union Statistics on Income and Living Conditions (EU-SILC). Panel B shows the fraction of households that, over the last year, were in arrears on their utility bill (PFV3). The fraction is given by the length of each bar. The vertical line indicates the average of the fractions in our sample of countries. Data are from European Union Statistics on Income and Living Conditions (EU-SILC).

country-specific aggregates, which relate to asset and income positions of households. Our primary focus is on the share of households living HtM, but we also report results for the four alternative measures introduced in Sections IIA and IIB: PFV1, Lottery MPC, PFV2, and PFV3.

Scatterplots between different measures of monetary policy effectiveness and the shares of households living HtM are presented in Figure 5. Both panels show the share of HtM households on their horizontal axes, and the vertical axes display different measures of the effectiveness of monetary policy. Figure 5, panel A shows the peak of the output impulse response, which exhibits a significant positive correlation with the HtM share across countries.²⁷ Figure 5, panel B instead uses the cumulative impulse response, with very similar results. Both suggest that an accommodating monetary policy shock has bigger effects on output in countries with a higher share of HtM households.

We interpret these results in light of a standard TANK model, as in Bilbiie (2020b). Here, a certain fraction of households consume their income every period, by construction, while the remainder can save and borrow. This simple setup captures an important element of monetary policy transmission with heterogeneous agents: a partial and a general equilibrium effect. The former describes output effects that occur due to the Euler equation of the unconstrained households. A shock that lowers the real interest rate makes these households demand more output in the current period. The general equilibrium effect includes the changes in output caused by changes in wages and profits. In the model, whether the share of constrained households amplifies (dampens) the aggregate output response depends on whether

²⁷ Because we calculate both the peak responses and the HtM shares, there is uncertainty associated with our point estimates. In order to not clutter the figure reported here, we relegate the scatterplot including confidence intervals to Figure A1, panel A in the Appendix.

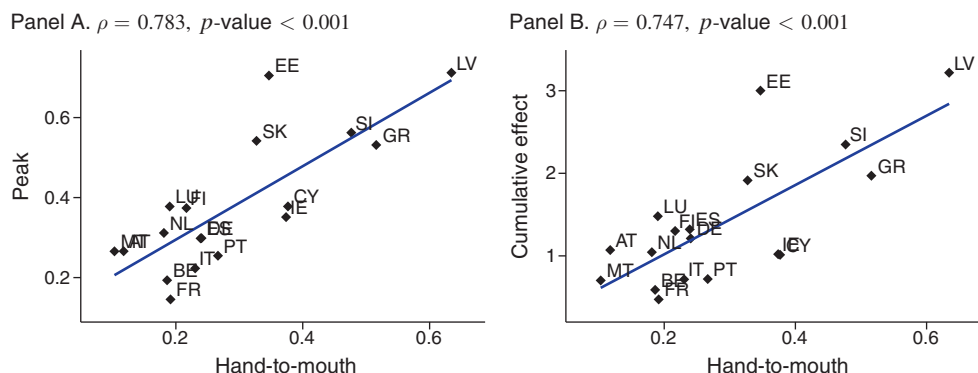


FIGURE 5. MONETARY POLICY EFFECTIVENESS AND HAND-TO-MOUTH SHARES

Notes: This figure plots the effectiveness of monetary policy as measured by the peak effect and cumulative effect of the real GDP impulse responses, calculated using the benchmark LPIV estimation, against the share of households classified as living HtM in each euro-area country (except Lithuania, not included in the HFCS). The HtM shares are calculated using data from the Eurosystem Household Finance and Consumption Survey. The impulse is an expansionary monetary policy shock of one standard deviation. The blue lines are fitted from regressions of peak/cumulative values on HtM shares. In the upper-left corner of each panel, we report the correlation coefficient ρ and the p -value. Panel A shows peak effects and share of HtM. Panel B shows cumulative effects and share of HtM, normalized by aggregate euro-area cumulative effect.

income—i.e., the sum of wage and profit income—of constrained households moves more (less) than one-for-one with aggregate income.

The results in Figure 5 show that a higher share of liquidity-constrained households amplifies an economy's response to monetary policy surprises. As explained above, this is in line with Bilbiie (2020b) if the income elasticity of constrained households with regard to aggregate income is larger than one. Richer models, such as those in Auclert (2019) or Hagedorn et al. (2019), feature more channels through which different households can be differently affected by aggregate shocks; still, they imply that if the income of the highest MPC agents co-moves more with aggregate income than that of the low-MPC agents, this mechanism amplifies the economy's response to shocks relative to RANK models.

Using the panel dimension of the HFCS, we find suggestive evidence that the elasticity of HtM households' three-year income growth with respect to aggregate (three-year) income growth is significantly higher than that of non-HtM households. For details, see Appendix A.A1. These findings are in line with Patterson (2022), who, in a US context, provides evidence for MPC being larger for individuals who are more affected by business cycles.

We now turn to our alternative measures of MPC, namely, PFV1–PFV3 and the self-reported propensity to consume out of lottery winnings. Our focus is on peak responses, but as before, results are similar using cumulative responses as the measure for monetary policy effectiveness.

The four scatterplots are presented in Figure 6. Correlations between the peak responses and each of the four statistics are strong. We view this as encouraging for two reasons. First, the results lend credence to the measure proposed by Kaplan,

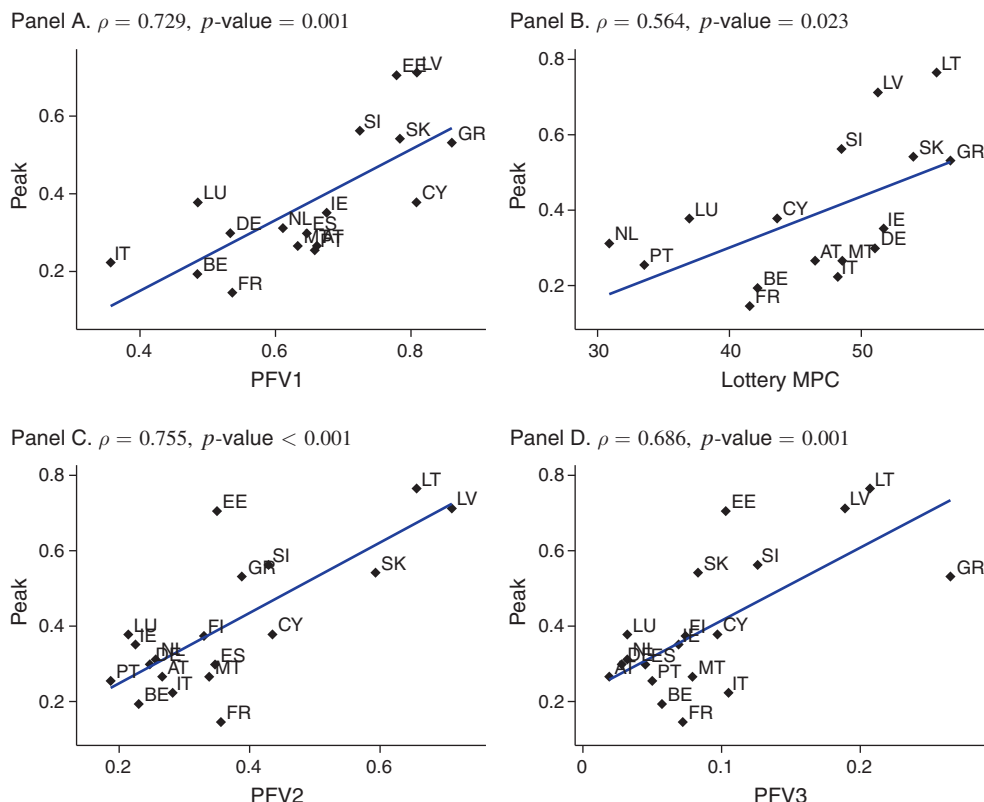


FIGURE 6. IMPACT OF MONETARY POLICY AND ALTERNATIVE LIQUIDITY CONSTRAINT MEASURES

Notes: This figure plots the effectiveness of monetary policy as measured by the peak effect, calculated using the benchmark LPIV estimation, against different statistics in each euro-area country. The impulse is an expansionary monetary policy shock of one standard deviation. The blue lines are fitted from regressions of peak values on the respective statistics. In the upper-left corner of each panel, we report the correlation coefficient ρ and the p -value. Panel A shows peak effects and PFV1. Panel B shows peak effects and lottery MPC, calculated using data from the Eurosystem Household Finance and Consumption Survey (HFCS). Panel C shows peak effects and PFV2. Panel D shows peak effects and PFV3. PFV1–PFV3 and lottery MPC are as defined in Section II. PFV1 is calculated using data from the HFCS, and to calculate PFV2 and PFV3, we use data from European Union Statistics on Income and Living Conditions (EU-SILC).

Violante, and Weidner (2014). The correlations are very similar, although the alternative proxies use different approaches and, in two cases, different surveys. Second, the proxies for MPC in panels C and D were calculated using data from 2005, giving us confidence that our results are not driven by the financial crisis or the European sovereign debt crisis. On the contrary, the correlations we find are a persistent feature of European monetary policy transmission.

Next, we investigate the importance of the distinction between liquid and illiquid asset holdings in more detail. Kaplan and Violante (2014) and Kaplan, Violante, and Weidner (2014) argue that it is important to disaggregate these two types of assets by partitioning households into wealthy HtM households (liquidity-constrained but owning positive illiquid wealth) and poor HtM households (zero or negative illiquid wealth). They estimate the MPC of P-HtM (W-HtM) households to be twice (thrice)

as large as the MPC of unconstrained households. However, a sole focus on differences between P-HtM and W-HtM households in MPC overlooks that their incomes might adjust differently and a potential revaluation of the illiquid asset portfolios of W-HtM households following a shock to the interest rate (Auclert 2019). Our data do not allow us to investigate how income and asset values change following the shocks.

Figure 7 shows the relationship between the peak responses across countries and their shares of wealthy and poor HtM households in panels A and B, respectively. While the W-HtM share is strongly correlated with the peak values of the IRFs, this is not the case for the share of P-HtM households. This suggests that disregarding households' liquidity positions in theoretical models and empirical work can lead to erroneous conclusions about the effects of monetary policy, as argued by Kaplan, Violante, and Weidner (2014). We view this as an interesting question for future research.

As a complementary test, we investigate the relationship between peak responses and the fraction of asset poor. Kaplan, Violante, and Weidner (2014) argue that total net wealth, which is the standard metric for high-MPC behavior in heterogeneous-agent macroeconomic models, is a poor predictor of MPC. As in Kaplan, Violante, and Weidner (2014), a household is labeled as asset poor if the sum of its net wealth is zero or negative. Figure 7, panel C gives no evidence for a relationship between output responses and the share of asset-poor households, and the statistic is outperformed by all of our other measures in predicting by how much output is affected through monetary policy shocks.

B. Robustness

First, we test whether our results are affected by restricting the sample to the countries who adopted the euro by the year 2002, when the currency was introduced. For this set of countries, the ECB was the relevant monetary policy institution throughout our sample period. The first column of Table 1, for reference, reports the results outlined in the previous section (reported in Figure 5, panel A). The second column reports the same statistics for the sample of initial euro-area members. Although the correlation between the share of HtM households and the peaks of the IRFs is attenuated slightly, it is still 0.7 and statistically significant. We see this as encouraging, as the conclusions drawn in the previous section are not driven by countries that joined the currency union after 2002.

Along similar lines, we can test whether using the first wave of the HFCS, conducted in 2010, affects our conclusions. Table 1, column 3 reports the correlation between HtM shares computed from the HFCS's first wave and peak responses of GDP after a monetary policy surprise. Importantly, during the first wave, the HFCS was not conducted in all countries in our sample, which leads us to restrict the analysis to the initial members of the euro area. The relevant comparison is the second column in Table 1. While the *t*-statistic becomes slightly smaller, the point estimates are almost equivalent across different survey waves.²⁸

²⁸ We report the fractions of HtM households across countries according to all three survey waves of the HFCS in Figure A4 in the Appendix. The fractions are remarkably stable across time.

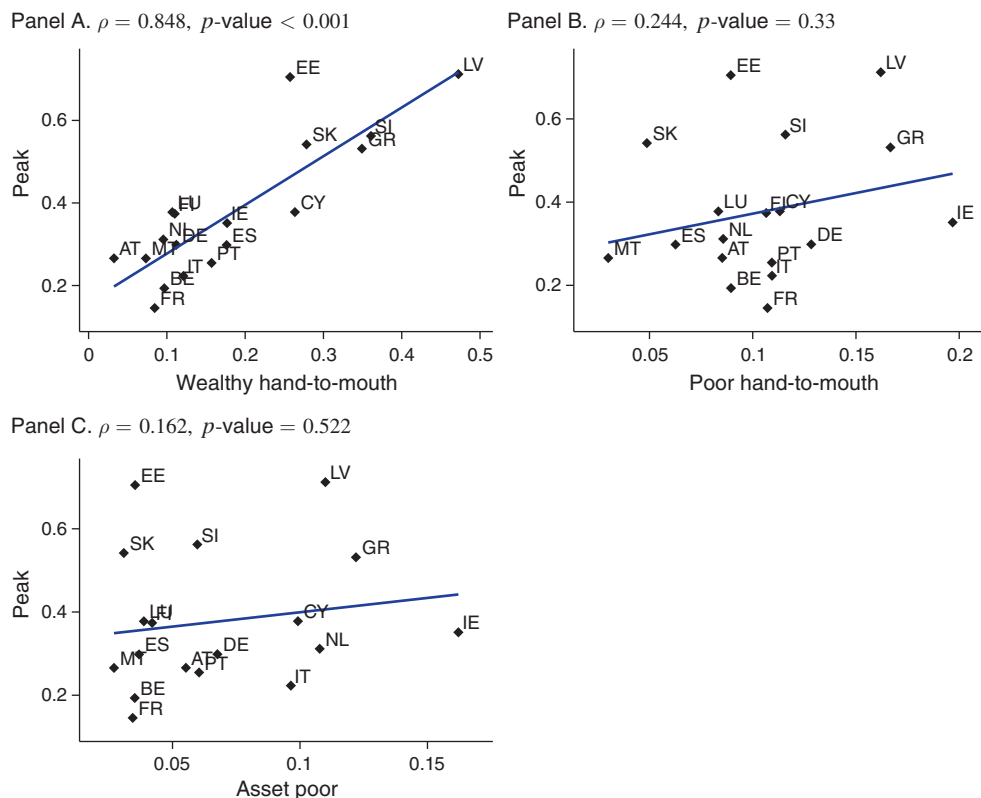


FIGURE 7. MONETARY POLICY EFFECTIVENESS AND WEALTHY AND POOR HAND-TO-MOUTH SHARES

Notes: This figure plots the effectiveness of monetary policy as measured by the peak effect, calculated using the benchmark LPIV estimation, against the share of households classified as living as wealthy HtM, poor HtM, and asset-poor, respectively, by Kaplan, Violante, and Weidner (2014) in each euro-area country (except Lithuania). The impulse is an expansionary monetary policy shock of one standard deviation. The blue lines are fitted from regressions of peak values on wealthy HtM shares, poor HtM shares, and the share of asset poor, respectively. In the upper-left corner of each panel we report the correlation coefficient ρ and the p -value. Panel A shows peak effects and share of wealthy HtM. Panel B shows peak effects and share of poor HtM. Panel C shows peak effects and the share of asset poor. Wealthy HtM shares, poor HtM shares and shares of asset poor are calculated using data from the Eurosystem Household Finance and Consumption Survey.

Because consumption as opposed to GDP is the relevant metric for household welfare, columns 4 and 5 in Table 1 repeat the exercise from Section IIIA, substituting a quarterly measure of household consumption for GDP. As before, we interpolate it to monthly frequency.²⁹ The results for the full sample (column 3) are very similar to those estimated using the monthly GDP series. The correlation coefficient falls very slightly from 0.78 to 0.75.³⁰ Restricting the sample to the initial members of the euro area (column 5), the correlation coefficient increases considerably to 0.88. The scatterplots associated with these estimations are reported in Figure 8. These

²⁹For each euro-area country we again use monthly data for industrial production, retail, trade, and unemployment to construct series for monthly household consumption.

³⁰See Figure A1, panel B for the associated scatterplot, including confidence bands.

TABLE 1—ROBUSTNESS TESTS

	Baseline	Init. members	1st wave	Consumption	Cons-IM	GVAR	JK
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ρ	0.78	0.70	0.69	0.75	0.88	0.78	0.58
t -statistic	5.04	3.07	2.87	4.43	5.84	4.91	2.86
p -value	0.00	0.01	0.02	0.00	0.01	0.00	0.01

Notes: The table shows the correlation coefficient ρ between a measure of the share of HtM households and the peak output response to a monetary policy shock across countries. The second and third row display the associated t -statistic and p -value. The first column shows the results for our baseline specification. The second column shows the results when we restrict the sample to euro-area countries that were members in 2002. The third column restricts the sample to early euro-area members using the first wave of the HFCS to compute HtM shares. The fourth column uses the peak response from a consumption IRF, as opposed to GDP. The fifth column uses peak responses from a consumption IRF and restricts the sample to early euro-area countries. The sixth column obtains the peak responses from a GVAR outlined in online Appendix B.B1. The seventh and last column uses peak responses produced using the monetary policy shocks series in Jarociński and Karadi (2020).

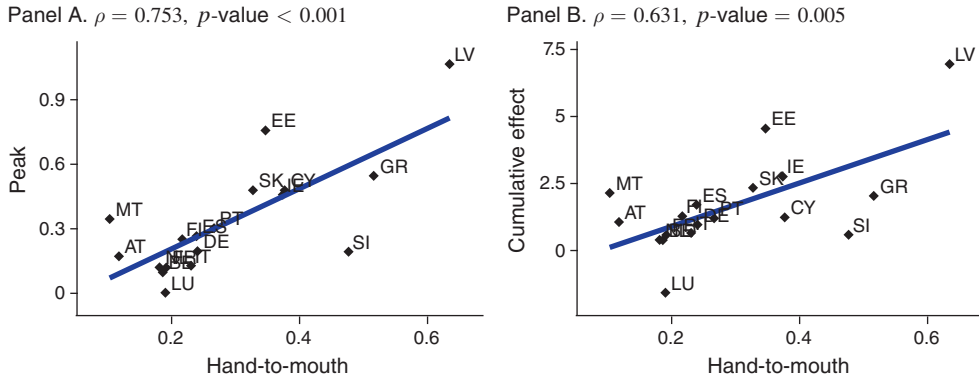


FIGURE 8. MONETARY POLICY EFFECTIVENESS AND HAND-TO-MOUTH SHARES—CONSUMPTION RESPONSES

Notes: This figure plots the effectiveness of monetary policy as measured by the peak effect and cumulative effect of the total household consumption impulse responses, calculated using the benchmark LPIV estimation, against the share of households classified as living HtM in each euro-area country (except Lithuania, not included in the HFCS). The HtM shares are calculated using data from the Eurosystem Household Finance and Consumption Survey. The impulse is an expansionary monetary policy shock of one standard deviation. The blue lines are fitted from regressions of peak/cumulative values on HtM shares. In the upper-left corner of each panel, we report the correlation coefficient ρ and the p -value. Panel A shows peak effects and share of HtM. Panel B shows cumulative effects and share of HtM.

results indicate that our initial findings are not driven by heterogeneous investment demand or fiscal responses across countries.³¹

As another robustness check, we construct a GVAR for the euro area based on the work by Georgiadis (2015) and Burriel and Galesi (2018) and repeat our analysis in this framework. (See online Appendix B.B1 for details on the model setup.) We utilize the same data as for the LPIV estimation and, to our knowledge, we are the first

³¹This conclusion assumes $GDP = C + I + G$.

to combine the instrumental VAR techniques laid out in Stock and Watson (2018) with the GVAR setting. The correlation coefficient between the peak responses estimated using this approach and the HtM shares across countries is reported in Table 1, column 6. It is the same as the same statistic obtained from the LPIV estimation and highly significant.

Lastly, we show that our results are robust to using the shock series produced by Jarociński and Karadi (2020), who distinguish between monetary policy shocks and information shocks. We use the former series and repeat the analysis above, constructing new impulse response functions and obtaining new peak values.³² Table 1, column 7 shows that the correlation statistic between the share of HtM households and the peak responses is lower than it is with our shock series but still highly significant.

Next, we show that our results are robust to changing the horizon at which the effects of monetary policy are measured. In the section above, we mainly rely on peak values. Here we instead focus on the point estimates at different horizons $h = \{0, 1, \dots, H\}$ and first extract the point estimate for each country n , β_n^h , to correlate each of these with the HtM values.

The horizontal axis in Figure 9, panel A shows the horizon (h), and the vertical axis shows the correlation between the country-specific HtM measures and IRF point estimates. During the majority of the first year, the correlation is not significant. This is unsurprising, as monetary policy affects output with a lag. After a year, however, the correlation is statistically and economically significant until it dies out toward the end of our estimation horizon. The latter, again, is unsurprising, as Figure 2 indicates that the effect of a common monetary policy shock peters out after three years in most countries.

Second, we divide the countries into two groups based on HtM shares. Countries with HtM shares below the median are placed in the first group, and countries with HtM shares above the median are placed in the second group. We then re-estimate equation (1) for each of the two groups.³³

Figure 9, panel B graphs the results for the two group-specific IRFs. We find again that output reacts more to monetary policy shocks in countries with higher HtM shares. We view these results as strengthening our previous conclusion that the share of HtM households is a relevant statistic for the effectiveness of monetary policy across countries.³⁴

Next, we investigate whether other country-specific characteristics can account for the heterogeneity in impulse responses that we observe. We focus on a set of variables that could be correlated with both HtM shares and the effectiveness of monetary policy. Our strategy is the following: First, we gather data on variables suggested in the literature as relevant for the effectiveness of monetary policy in a

³² The resulting impulse responses for output and prices are reported in Appendix A.A2.

³³ After dividing countries into the two groups, we then index the GDP series of each country and use the average index values within each group as a measure for GDP.

³⁴ We can perform the same analysis using a Panel IV setup, including country fixed effects. This approach is discussed in online Appendix B.B2, where we show that the inclusion of such fixed effects does not change our conclusions.

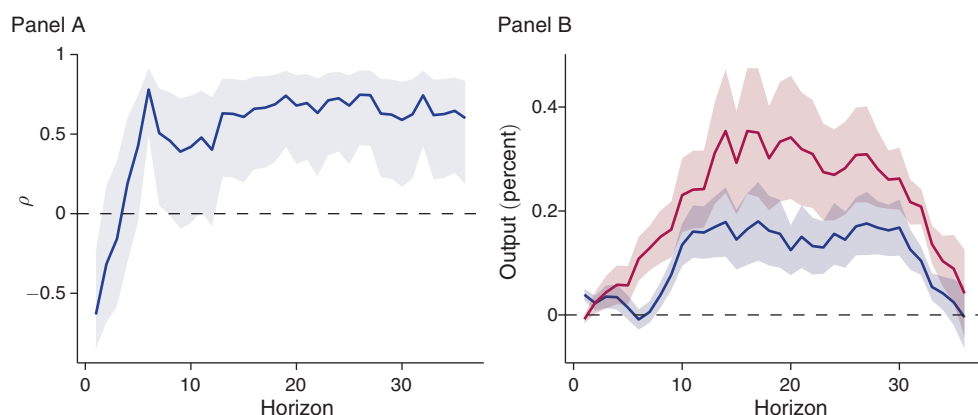


FIGURE 9. MONETARY POLICY EFFECTIVENESS AND WEALTHY AND POOR HAND-TO-MOUTH SHARES

Notes: Panel A shows correlation between HtM and responses at different horizons, retrieved from the impulse responses using the benchmark specification, for each horizon $h = \{0, 1, \dots, H\}$. The shaded area is the 95 percent confidence band around the point estimates. Panel B shows IRFs for two groups of countries. The blue line represents the IRF for the group consisting of countries with HtM shares below the median, and the red line represents the IRF for the group consisting of countries with HtM shares above the median. (See text for details.) The shaded areas give 68 percent confidence bands around the point estimates. Calculations of HtM shares are based on data from the Eurosystem Household Finance and Consumption Survey.

cross-country perspective. Subsequently, for each variable we investigate whether it is correlated with HtM shares, whether it is correlated with the peak effects we find in Section I, and whether after controlling for the variable, the HtM share still explains a part of the output responses we observe.³⁵

Cloyne, Ferreira, and Surico (2020) find that households that own a mortgaged property adjust consumption spending more than both renters and homeowners without mortgages in response to unexpected interest rate changes. The authors find that consumption among homeowners without mortgages is insensitive to changes in monetary policy. It is, furthermore, possible that monetary policy can affect housing prices and output via the collateral channel (Cloyne et al. 2019). Corsetti, Duarte, and Mann (2018) find that the strength of the housing channel is related to home ownership rates. These results lead us to the first three variables that are introduced in this section. The first variable is labelled *Own* and represents the fraction of households in each country that own their main residence. We allow for an outstanding mortgage to be tied to the residence. A closely related variable is *Mort*, which represents the fraction of households in each country that have a mortgage. Additionally, in each country there are households that own their main residences but do not have mortgages attached to them. We label the variable for the fraction of these households in each country as *HO*.

It is possible that the effectiveness of monetary policy depends on how highly indebted households are (see, e.g., Flodén et al. 2021) and on how common it is that

³⁵ Most of these control variables are constructed using data from the HFCS, since many of them are related to housing and how it is financed.

mortgages have an adjustable interest rate (see, e.g., Calza, Monacelli, and Stracca 2013; Flodén et al. 2021). We calculate the fraction of households that have at least one mortgage with an adjustable interest rate and label the variable *Flex*. To test whether HtM shares and effectiveness of monetary policy are related to how highly indebted households are, we calculate average loan-to-value ratios and average loan-to-income ratios among households with mortgages in each country and label them *LTV* and *LTI*, respectively. Observations (households) with *LTV* above 1.5 were removed in the calculations of *LTV* and observations with *LTI* above 10 were removed from the calculations of *LTI* to limit the influence of outliers.

In Section IIIA we argue that it is mainly the fraction of wealthy HtM households that explains why the total fraction of HtM households is correlated with peak values. It is possible that this result is driven by the share of households with positive amounts of illiquid wealth, not necessarily by the share of HtM. We can rule this out by showing that there is much variation in wealthy HtM shares that is not due to the variation in the shares of wealthy people that is correlated with peak values. We therefore calculate the share of wealthy households in each country and label the variable *Wealthy*.³⁶

Wong (2021) finds that in the United States, especially, younger households refinance loans following changes in the interest rate and drive most of the aggregate response in consumption. Examining the HFCS data, we observe that older households, on average, are less likely to be HtM. Moreover, the probability of being wealthy HtM increases between age 20 and the late 30s and decreases after this threshold. We find it important to test whether including the average age in each country as a control variable changes our results. The average age of household heads in our data is calculated and is labelled *Age*.

The growing literature using GVAR models (e.g., Burriel and Galesi 2018; Georgiadis 2015) emphasizes the importance of considering spillover effects of monetary policy and whether the size of these spillovers is partly related to trade flows. We include a measure of trade openness due to its importance in the Dynamic IS equation in the small open economy literature (Galí and Monacelli 2008). We calculate trade openness as the sum of imports and exports as a share of GDP in each country to test whether what we find is related to trade. We use the World Bank national accounts data to calculate this statistic and label it *Trade*.

The next variable is labelled *ROL* and is related to how regulated labor markets are. Georgiadis (2015), using data from a subset of the countries that we consider, estimates that output in countries with more regulation responds less to monetary policy shocks. We construct it by calculating the average of the “Employment laws index” and the “Collective relations laws index” from Botero et al. (2004). Georgiadis (2015) also finds that the share of GDP accounted for by services is closely connected to the effectiveness to monetary policy, showing that countries that have the lowest shares compared to countries with the highest shares exhibit responses of output that are half as large. We have mentioned that our estimates for the effectiveness of monetary policy are similar to the estimates in his paper. Hence,

³⁶ A household whose net illiquid assets are positive is labeled as wealthy.

TABLE 2—CORRELATIONS AND SEMIPARTIAL CORRELATIONS

<i>X</i>	$\rho(\text{Peak}, X)$		$\rho(\text{HtM}, X)$		$\rho(\text{Peak}, \text{HtM} - X)$	
Peak	1.00	(0.000)	0.78	(0.000)	NA	(NA)
HtM	0.78	(0.000)	1.00	(0.000)	NA	(NA)
Own	0.42	(0.086)	0.36	(0.143)	0.68	(0.003)
Mort	−0.35	(0.155)	−0.32	(0.196)	0.71	(0.001)
HO	0.54	(0.022)	0.47	(0.047)	0.60	(0.011)
Wealthy	0.35	(0.148)	0.23	(0.35)	0.72	(0.001)
Flex	−0.04	(0.894)	−0.03	(0.914)	0.79	(0.000)
Age	−0.05	(0.851)	0.10	(0.703)	0.79	(0.000)
LTV	−0.34	(0.163)	−0.26	(0.296)	0.72	(0.001)
LTI	−0.37	(0.135)	−0.22	(0.37)	0.72	(0.001)
Trade	0.10	(0.67)	−0.18	(0.483)	0.81	(0.000)
ROL	−0.08	(0.797)	0.11	(0.726)	0.88	(0.000)
Services	−0.41	(0.083)	−0.19	(0.442)	0.73	(0.001)
HP growth	0.34	(0.166)	−0.13	(0.623)	0.83	(0.000)
GDPpc	−0.44	(0.058)	−0.47	(0.048)	0.68	(0.003)

Notes: The first column shows the correlation coefficient between estimated peak values and the variables that vary across the rows in the table. The second column shows the correlation coefficient between HtM shares and the variables that vary across the rows in the table. The third column shows the semipartial correlation between the estimated peak values and HtM shares. The *p*-values for the correlation coefficients are reported within parentheses to the right of each coefficient. Calculations of *Peak*, *HtM*, *Own*, *Mort*, *HO*, *Wealthy*, *Flex*, *Age*, *LTV*, and *LTI* are based on data from the HFCS. *HP Growth* is the average quarterly year-on-year growth in Eurostat's housing price index from the first data point (2006:I for most countries) until 2012:IV. Greece is missing from the index. GDP per capita is for 2008. See text for information about the sources of the other variables.

it is likely that our measures are also correlated with service shares, and it becomes important to see if there is variation left in HtM shares even after having controlled for service shares that is correlated with effectiveness of monetary policy. We label the variable *Service* and to calculate it, we average over the shares reported in the World Bank's WDI database between the years 2000 and 2012 for each country.

Economic development is potentially correlated with how countries respond to shocks and with the share of HtM households. For this reason, we control for GDP per capita of 2008.

Our sample period coincides with large fluctuations in housing prices in some European countries. In order to show that the size of our HtM shares is uncorrelated with these changes, we control for a measure of housing price growth across European countries. We utilize Eurostat's housing price index, which starts in 2005. Housing price growth is calculated as the average quarterly year-on-year change in the index between the first quarter for which data are available and the last quarter of 2012.³⁷

All results are summarized in Table 2. The first column presents raw correlations between the peak effects and the different variables that vary across the rows in the table. In the second column we see the correlations between the HtM shares and the variables that vary across the rows. Most often the absolute values of the correlation coefficients are relatively close to zero. One exception is *HO*, for which the correlation is positive and of significant magnitude with both peak effects and

³⁷ For most countries, the first data point is available in 2005. Data for Italy and Austria are only available since 2010. The index is not available for Greece.

HtM shares.³⁸ Another is *Services*, which is negatively correlated with peak effects (confirming the result from Georgiadis 2015) and also negatively correlated with HtM shares.

That peak effects and/or HtM shares are correlated with some of these variables was expected. The important question is whether these other variables are likely to be the reason we find such a strong correlation between peak responses and HtM shares. To get a sense of whether this could be the case, we calculate semipartial correlations between the estimated peak effects and HtM shares. These semipartial correlations are reported in the third column of Table 2 and indicate the correlation coefficient between the peak effects and HtM shares after the variation in HtM shares explained by these other variables, varying across the rows in the table, has been accounted for. Going down the rows, we conclude that the coefficient remains large. From being 0.78 without having controlled for any other variable, it reaches its lowest value at 0.60 when we account for the variation in *HtM* explained by *HO* and as high as 0.88 when we instead extract the variation in *HtM* explained by *ROL*. Based on the results presented in Table 2, we find no variable that supports the conclusion that the correlation between the peak effects and HtM shares is driven by omitted variables. For example, it could have been the case that all HtM households but no non-HtM households had mortgages. In such a case, the correlation between output responses and shares of HtM could potentially be explained by the fact that higher shares of households with mortgages caused larger output responses. The results presented in Table 2 suggest that a higher fraction of constrained households causes output to respond more to monetary policy shocks.

Intuitively, many of the variables considered in the table, such as home ownership (*Own*) and mortgage holdings (*Mort*), seem closely related to the HtM status of a household. Hence, it may be surprising that none of the variables in the table are able to attenuate the correlation we find significantly. It is important to realize, however, that none of the variables in Table 2, except for the constructed variable *HtM* itself, take the liquidity of a household's asset positions into account. In particular, the latter quantifies the relationship *liquid-assets-to-income*. Together with the estimated effects for monetary policy shocks on output, the results in Table 2 suggest that, if one is to construct a statistic based on household asset data with the intent to capture MPC, then no single variable is satisfactory, but one must classify assets based on liquidity and set them in relation to income.

IV. Conclusion

The introduction of heterogeneous agents into New Keynesian models is becoming widespread. However, there is still a lack of empirical evidence on how household heterogeneity in income and wealth affects the response of aggregate output following a monetary policy shock. In this paper we provide such evidence, showing that

³⁸ The negative correlation between *Mort* and *HtM* might seem surprising since it appears plausible that wealthy HtM households often have mortgages. In online Appendix B.B5 we show this to not be the case. Another potentially surprising finding, given results in Cloyne, Ferreira, and Surico (2020), is the positive correlation between *Peak* and *HO*. We investigate and discuss it further in online Appendix B.B5.

aggregate output responses are larger in countries with a higher share of liquidity-constrained households.

We estimate country-specific output responses in the euro area following an expansionary monetary policy shock. The IRFs are produced using Local Projections (Jordà 2005). To identify surprise changes in the policy rate, we construct an instrument based on movements in Eonia OIS rates during a narrow time window around the ECB's monetary policy announcement and the subsequent press conference. Given that the countries within the euro area share a central bank, we can rule out that any heterogeneity in IRFs is due to differences in the success of our identification method across countries.

We find that output responses to a common monetary policy shock in the euro area are heterogeneous across countries in terms of cumulative impact and peak values.

Subsequently, we correlate the country-specific responses with proxies for the share of liquidity-constrained households across countries. Intuitively, these households are less able to smooth income fluctuations following monetary policy shocks. Our main measure is the share of households that are classified as HtM, according to the definition by Kaplan, Violante, and Weidner (2014), but we construct four additional measures of the share of constrained households, which are distinct in the surveys and time periods used to construct them.

On average, countries with a higher share of liquidity-constrained households react more strongly to a monetary policy shock. When splitting the sample by shares of HtM households, the aggregate response of the high-HtM countries is significantly stronger than that of the low-HtM countries. These findings are in line with theoretical work, given plausible assumptions about the elasticity of constrained households' incomes to aggregate income (Bilbiie 2020b).

Our findings support the notion that research on monetary policy needs to account for heterogeneity across the income and wealth distributions. Furthermore, they imply that liquidity is an important factor in how monetary policy shocks affect households and the real economy. Additional empirical research is needed, however, to understand the mechanism through which this heterogeneity in liquidity directly shapes the responses of output to monetary policy shocks. We consider this a fruitful avenue for future research.

APPENDIX A

A1. *Income Elasticities*

The amplification result outlined in Bilbiie (2020a) requires that constrained (unconstrained) households' income elasticities with respect to aggregate income be larger (smaller) than one. Empirical evidence to this effect is scarce.³⁹ We therefore test for the income elasticity mechanism using the HFCS dataset.

³⁹Coibion et al. (2017) find that inequality rises after contractionary monetary policy in the United States. They estimate that the change in labor earnings of high-net-worth households is lower than that of low-net-worth households after monetary shocks and that incomes of households at the ninetieth percentile rise somewhat relative to the median household, while households at the tenth percentile see their relative incomes fall particularly sharply.

A subset of households in our sample, from a subset of countries that participate in the HFCS, are interviewed in multiple survey waves. We use data for these households and investigate their income elasticities with respect to aggregate income. Since data from three waves currently exist, we compute the individual growth rates between the first and second waves and between the second and third waves. To limit the influence of outliers, households whose incomes or income growth rates were below or above the first and ninety-ninth percentiles, respectively, in each country and time period were removed.⁴⁰ Since the HtM status of a household can change between the survey waves, we choose to classify a household as HtM if it was classified thusly in the first wave contributing to the income growth rate.⁴¹ Sample weights are employed in the estimation.

We run the following regression, following, e.g., Guvenen, Ozkan, and Song (2014), but distinguishing by HtM status:

$$(A1) \quad \Delta y_{i,n,t} = \frac{\alpha}{\begin{matrix} [9.307] \\ (1.143) \end{matrix}} + \frac{\beta}{\begin{matrix} [-1.401] \\ (2.670) \end{matrix}} HtM_{i,n,t-1} + \frac{\gamma}{\begin{matrix} [1.168] \\ (0.137) \end{matrix}} \Delta Y_{n,t} + \frac{\delta}{\begin{matrix} [0.643] \\ (0.326) \end{matrix}} \Delta Y_{n,t} \\ \times HtM_{i,n,t-1} + e_{i,n,t}$$

where the left-hand-side variable is the growth rate of labor income for household i in country n between two periods $t - 1$ and t , HtM is the variable that indicates the Hand-to-Mouth status of the household (in period $t - 1$), and $\Delta Y_{n,t}$ is the growth rate of aggregate income in country n between periods $t - 1$ and t . Lastly, the regression includes an interaction between aggregate income growth and Hand-to-Mouth status. The coefficients of interest are γ and δ , where γ captures the (average) elasticity of individual income growth with respect to aggregate income growth for unconstrained households, and $\gamma + \delta$ captures the (average) elasticity of individual income growth with respect to average income growth for financially constrained households.

The estimated coefficients are reported below their respective parameters, and standard errors are placed inside parentheses.⁴² The first coefficient of interest, γ , is estimated to be 1.17 and is statistically significant at the 95 percent level.⁴³ On the other hand, δ is estimated to be 0.64 and is statistically significant at the 95 percent level (p -value 0.049). The value indicates that a 1 percentage point increase in

Patterson (2022) documents a positive covariance between workers' MPC and their earnings elasticity to GDP that is large enough to increase shock amplification.

⁴⁰The result presented in equation (A1) is robust to trimming below and above the fifth and ninetieth percentiles, respectively.

⁴¹As is discussed in more detail in online Appendix B.B4, the HFCS imputes data for missing values for some variables, and this is done five times, which results in five implicates. As a result of the imputation, the HtM status that we assign to households can vary across implicates. For the exercise that we perform in the current section, we classify a household as HtM if it was classified as HtM in at least three out of five implicates.

⁴²Robust standard errors are clustered at the individual level. We explore other alternatives, like country-level clustered standard errors, or estimate the standard errors using a wild bootstrap with standard errors clustered at the country level. The former yields a δ coefficient that is statistically significant at the 95 percent confidence level with only 12 clusters and the δ coefficient is statistically significant at the 90 percent level in the latter case.

⁴³Within each country, higher levels of income are associated with lower levels of income growth. It has the consequence that average income growth exceeds aggregate income growth, which explains why the estimated value for γ is greater than one.

aggregate income is associated with financially constrained households' incomes increasing by 0.64 percentage points more than those of unconstrained agents. Taken together, these findings suggest that if aggregate income grows, the income of financially constrained households grows by more and would, through the lens of Bilbiie (2020a), lead to amplification, as our results in Section IIIA suggest.

A2. Additional Figures and Tables

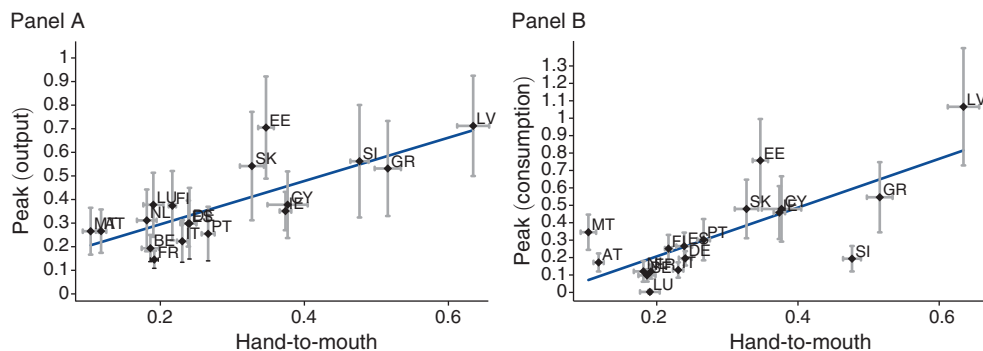


FIGURE A1. MONETARY POLICY EFFECTIVENESS AND HAND-TO-MOUTH SHARES—OUTPUT AND CONSUMPTION RESPONSES WITH CONFIDENCE BANDS

Notes: This figure plots HtM shares against peak responses of output (panel A) and consumption (panel B). The HtM shares are calculated using data from the Eurosystem Household Finance and Consumption Survey. The vertical lines and horizontal lines represent one-standard-deviation confidence bands for the peak responses and HtM shares, respectively. See online Appendix B.4 for more information about the standard errors for HtM.

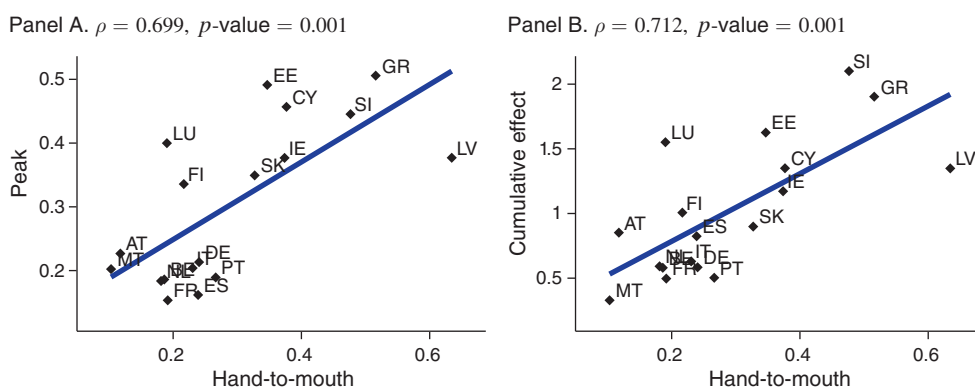


FIGURE A2. ROBUSTNESS INCLUDING COUNTRY-SPECIFIC LAGS

Notes: This figure plots the effectiveness of monetary policy as measured by the peak effect and cumulative effect of the real GDP impulse responses, calculated using the LPIV estimation with country-specific lags, against the share of households classified as living HtM in each euro-area country (except Lithuania, which is not included in the HFCS). The LPIV estimation includes three country-specific lags. The HtM shares are calculated using data from the Eurosystem Household Finance and Consumption Survey. The impulse is an expansionary monetary policy shock of one standard deviation. The blue lines are fitted from regressions of peak/cumulative values on HtM shares. In the upper-left corner of each panel, we report the correlation coefficient ρ and the p -value. Panel A shows peak effects and share of HtM. Panel B shows cumulative effects and share of HtM, normalized by aggregate euro-area cumulative effect.

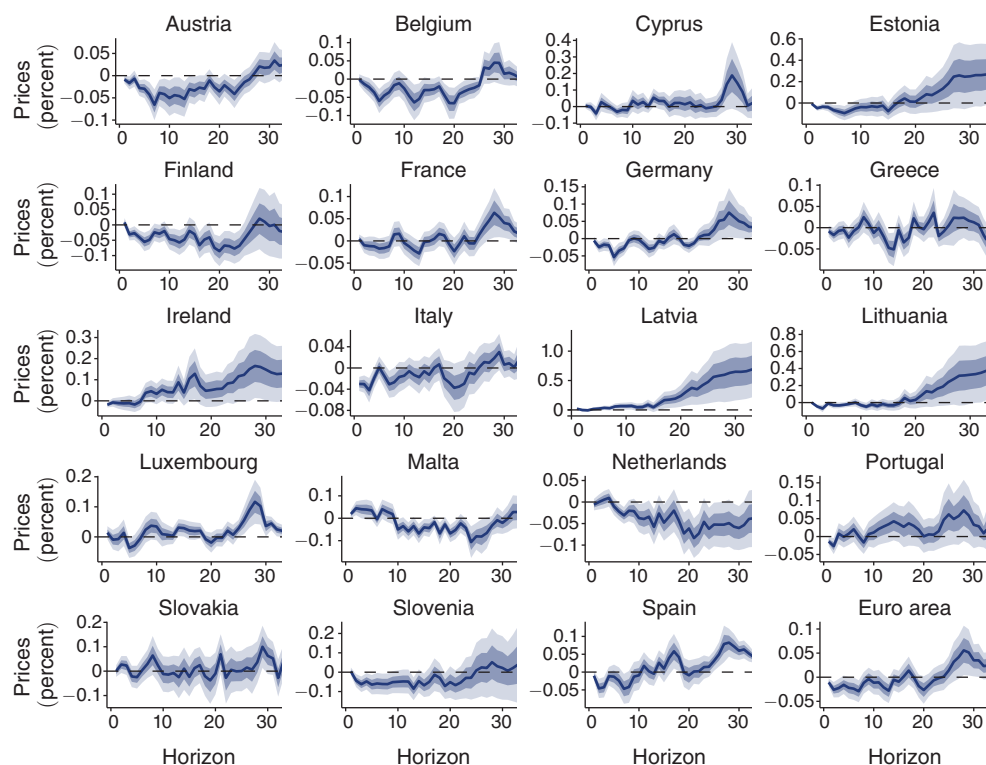


FIGURE A3. IMPULSE RESPONSES FOR PRICES IN EURO-AREA COUNTRIES—LPIV

Notes: This figure shows impulse responses of prices to an expansionary monetary policy shock of one standard deviation. For each euro-area country, the response is estimated using LPIV (equation (1)). The solid blue lines represent the IRFs produced by our preferred specification (see text for details). The dark blue and light blue shaded areas represent one- and two-standard-deviation confidence bands, constructed using Newey-West estimators. Note that the y-axes are scaled differently across countries.

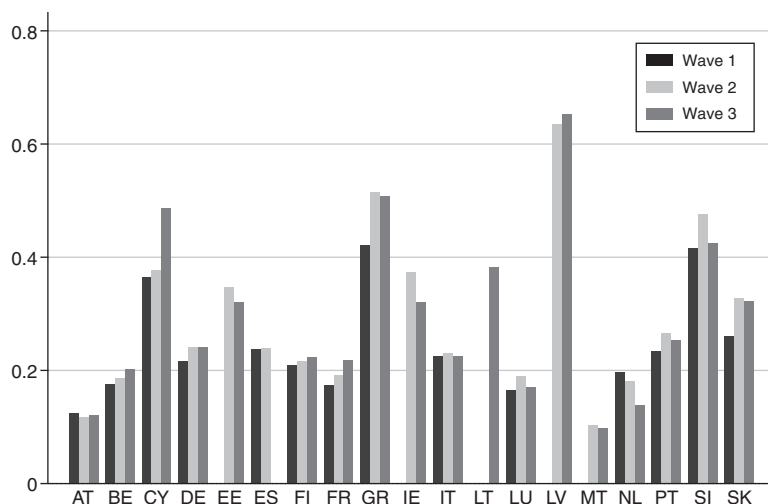


FIGURE A4. FRACTION OF HtM HOUSEHOLDS ACROSS HFCS SURVEY WAVES

Notes: This figure shows the fraction of HtM households, calculated according to the approach in Kaplan, Violante, and Weidner (2014), utilizing data from three different survey waves of the Eurosystem Household Finance and Consumption Survey.

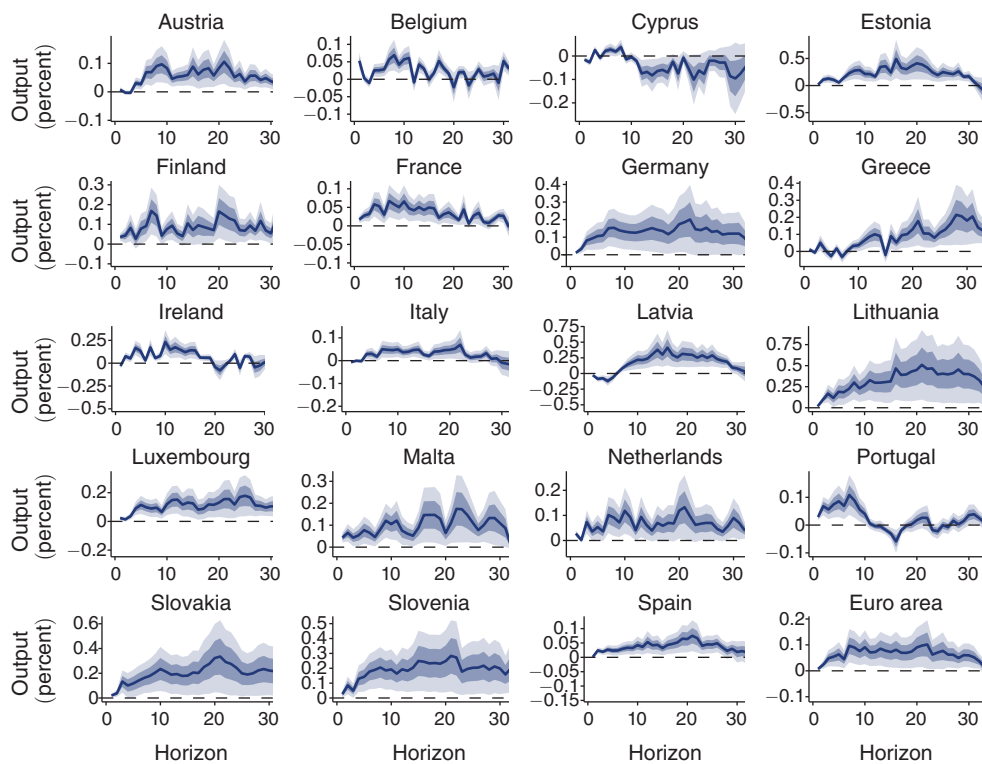


FIGURE A5. IMPULSE RESPONSES FOR OUTPUT IN EURO-AREA COUNTRIES—LPIV (JK)

Notes: This figure shows impulse responses of real GDP to an expansionary monetary policy shock of one standard deviation. The shock series is the monetary policy shock series reported in Jarociński and Karadi (2020). For each euro-area country, the response is estimated using LPIV (equation (1)). The solid blue lines represent the IRFs produced by our preferred specification (see text for details). The dark blue and light blue shaded areas represent one- and two-standard-deviation confidence bands, constructed using Newey-West estimators. Note that the y-axes are scaled differently across countries.

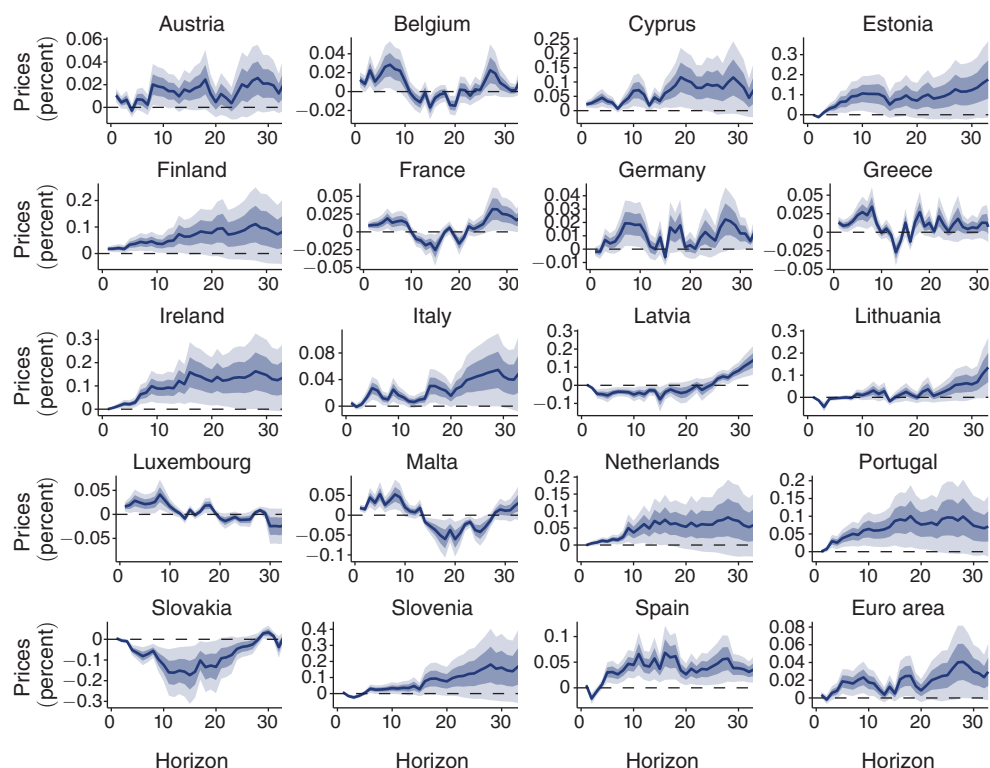


FIGURE A6. IMPULSE RESPONSES FOR PRICES IN EURO-AREA COUNTRIES—LPIV (JK)

Notes: This figure shows impulse responses of prices to an expansionary monetary policy shock of one standard deviation. The shock series is the Monetary Policy shock series reported in Jarociński and Karadi (2020). For each euro area country, the response is estimated using LPIV (equation (1)). The solid blue lines represent the IRFs produced by our preferred specification (see text for details). The dark and light blue shaded areas represent 1 and 2 standard deviation confidence bands, constructed using Newey-West estimators. Note that the y-axes are scaled differently across countries.

REFERENCES

- Almgren, Mattias, José-Elías Gallegos, John Kramer, and Ricardo Lima. 2022. “Replication Data for: Monetary Policy and Liquidity Constraints: Evidence from the Euro Area.” American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E140201V1>.
- Altavilla, Carlo, Domenico Giannone, and Michele Lenza. 2016. “The Financial and Macroeconomic Effects of the OMT Announcements.” *International Journal of Central Banking* 12 (3): 29–57.
- Ampudia, Miguel, and Skander Van den Heuvel. 2018. “Monetary Policy and Bank Equity Values in a Time of Low Interest Rates.” European Central Bank Working Paper 2199.
- Auclert, Adrien. 2019. “Monetary Policy and the Redistribution Channel.” *American Economic Review* 109 (6): 2333–67.
- Bilbiie, Florin O. 2008. “Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic.” *Journal of Economic Theory* 140 (1): 162–96.
- Bilbiie, Florin O. 2020a. “Monetary Policy and Heterogeneity: An Analytical Framework.” Unpublished.
- Bilbiie, Florin O. 2020b. “The New Keynesian Cross.” *Journal of Monetary Economics* 114: 90–108.
- Botero, Juan C., Simeon Djankov, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer. 2004. “The Regulation of Labor.” *Quarterly Journal of Economics* 119 (4): 1339–82.

- Burriel, Pablo, and Alessandro Galesi. 2018. "Uncovering the Heterogeneous Effects of ECB Unconventional Monetary Policies Across Euro Area Countries." *European Economic Review* 101: 210–29.
- Calza, Alessandro, Tommaso Monacelli, and Livio Stracca. 2013. "Housing Finance and Monetary Policy." *Journal of the European Economic Association* 11 (S1): 101–22.
- Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. 2017. "The Distribution of Wealth and the Marginal Propensity to Consume." *Quantitative Economics* 8 (3): 977–1020.
- Chow, Gregory C., and An-loh Lin. 1971. "Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series." *Review of Economics and Statistics* 53 (4): 372–75.
- Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary Policy When Households have Debt: New Evidence on the Transmission Mechanism." *Review of Economic Studies* 87 (1): 102–29.
- Cloyne, James, Kilian Huber, Ethan Ilzetzki, and Henrik Kleven. 2019. "The Effect of House Prices on Household Borrowing: A New Approach." *American Economic Review* 109 (6): 2104–36.
- Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2017. "Innocent Bystanders? Monetary Policy and Inequality." *Journal of Monetary Economics* 88: 70–89.
- Corsetti, Giancarlo, Joao B. Duarte, and Samuel Mann. 2018. "One Money, Many Markets—A Factor Model Approach to Monetary Policy in the Euro Area with High-Frequency Identification." Unpublished.
- Fagereng, Andreas, Martin B. Holm, and Gisle J. James Natvik. 2021. "MPC Heterogeneity and Household Balance Sheets." *American Economic Journal: Macroeconomics* 13 (4): 1–54.
- Flodén, Martin, Matilda Kilström, Jósef Sigurdsson, and Roine Vestman. 2021. "Household Debt and Monetary Policy: Revealing the Cash-Flow Channel." *Economic Journal* 131 (636): 1742–71.
- Galí, Jordi, J. David López-Salido, and Javier Vallés. 2007. "Understanding the Effects of Government Spending on Consumption." *Journal of the European Economic Association* 5 (1): 227–70.
- Galí, Jordi, and Tommaso Monacelli. 2008. "Optimal Monetary and Fiscal Policy in a Currency Union." *Journal of International Economics* 76 (1): 116–32.
- Georgiadis, Georgios. 2015. "Examining Asymmetries in the Transmission of Monetary Policy in the Euro Area: Evidence from a Mixed Cross-Section Global VAR Model." *European Economic Review* 75: 195–215.
- Gerlach, Stefan, and Frank Smets. 1995. "The Monetary Transmission Mechanism: Evidence from the G-7 Countries." Bank for International Settlements Working Paper 26.
- Gertler, Mark, and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7 (1): 44–76.
- Gornemann, Nils, Keith Kuester, and Makoto Nakajima. 2016. "Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy." Board of Governors of the Federal Reserve System International Finance Discussion Paper 1167.
- Güvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy* 122 (3): 621–60.
- Hagedorn, Marcus, Jinfeng Luo, Iouri Manovskii, and Kurt Mitman. 2019. "Forward Guidance." *Journal of Monetary Economics* 102: 1–23.
- Honkkila, Juha, and Ilja Kristian Kavonius. 2013. "Micro and Macro Analysis on Household Income, Wealth and Saving in the Euro Area." European Central Bank Working Paper 1619.
- 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): 161–82.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review* 108 (3): 697–743.
- Kaplan, Greg, and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82 (4): 1199–1239.
- Kaplan, Greg, Giovanni L. Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity* 45 (1): 77–138.
- Kuttner, Kenneth N. 2001. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics* 47 (3): 523–44.
- Lloyd, Simon P. 2018. "Estimating Nominal Interest Rate Expectations: Overnight Indexed Swaps and the Term Structure." Bank of England Staff Working Paper 763.

- Mandler, Martin, Michael Scharnagl, and Ute Volz.** 2016. "Heterogeneity in Euro-Area Monetary Policy Transmission: Results from a Large Multi-Country BVAR Model." Deutsche Bundesbank Discussion Paper 03/2016.
- Nakamura, Emi, and Jón Steinsson.** 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics* 133 (3): 1283–1330.
- Patterson, Christina.** 2022. "The Matching Multiplier and the Amplification of Recessions." Unpublished.
- Ramey, V.A.** 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics*, Vol. 2, edited by John B. Taylor and Harald Uhlig, 71–162. Amsterdam: Elsevier.
- Stock, James H., and Mark W. Watson.** 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *Economic Journal* 128 (610): 917–48.
- Werning, Iván.** 2015. "Incomplete Markets and Aggregate Demand." Unpublished.
- Wong, Arlene.** 2021. "Refinancing and the Transmission of Monetary Policy to Consumption." Unpublished.