

MONETARY POLICY AND LIQUIDITY CONSTRAINTS: EVIDENCE FROM THE EURO AREA*

Mattias Almgren
Jose E. Gallegos[†]
John Kramer
Ricardo Lima

September 18, 2019

Abstract

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 is larger in countries that have a larger fraction of households that are liquidity constrained.

JEL Classifications: E21, E31, E44, E52, F45.

*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 Adrien Auclert, Martín Beraja, Xavier Gabaix, Alessandro Galesi, Pierre-Olivier Gourinchas, John Hassler, Greg Kaplan, Peter Karadi, Karin Kinnerud, Kasper Kragh-Sørensen, Jesper Lindé, Haroon Mumtaz, Morten Ravn, Giovanni Ricco, Federica Romei, Maria Sandström, Jósef Sigurdsson, Xueping Sun, Javier Vallés, Has van Vlokhoven, Iván Werning and seminar participants at the IIES, Stanford University, Stockholm University and the Stockholm School of Economics for useful feedback and comments. Björn Hagströmer generously aided us in accessing the swap-rate data. Jose E. Gallegos and John Kramer greatly acknowledge funding from the Tom Hedelius foundation. 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 results or analysis of the data producers.

[†]Corresponding author: jose.elias.gallegos@iies.su.se, +34 650 662 152, Institute for International Economic Studies (IIES), Stockholm University, SE-106 91, Stockholm, Sweden.

I INTRODUCTION

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

While theoretical research has shown that heterogeneous agent models which include constrained agents can have different policy implications than their representative agent counterparts (Hagedorn et al. 2019), empirical evidence on how heterogeneity matters for the transmission from monetary policy to output is scant. 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 country 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. Since standard VAR techniques are infeasible in a system of many countries that share a policy rate, we rely on a Local Projection (LP) approach pioneered by Jordà (2005). This makes it possible to transparently estimate the same equation for all countries in our sample. 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). 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 sur-

1. According to the European Union Survey of Income and Living conditions. The monetary values represent the country-specific at-risk-of-poverty threshold, defined as 60 % of the national median equivalized disposable income after social transfers.

prises. Because OIS 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. Our 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 leads 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 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. In it, participating households are 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 output responses to a common European monetary policy surprise are not homogeneous across countries. There is significant heterogeneity in terms of cumulative impact, peak values and timing. 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 a stronger cumulative output response and a bigger peak response 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, we show that it is possible to estimate IRFs for countries within a currency union transparently, relying on

LP methods. Second, the finding that IRFs to a common monetary policy shock are heterogeneous across the euro area implies that a focus on European aggregates may be insufficient for both policymakers and researchers. Understanding the reasons for the differences we uncover is crucial in order to calibrate future policies. Lastly, we show that heterogeneity in general, and household balance sheets in particular, affect the transmission of monetary policy to the economy. 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.

Our research is related to several strands of literature. There is a large body of research which performs cross-country monetary policy analysis. An early example is Gerlach and Smets (1995) who perform a Structural VAR analysis of the G-7 countries and find that responses to country-specific monetary policy shocks are similar. 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 (2018a) 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 U.S. In the European context, there are no financial instruments equivalent to Fed funds futures which has led researchers to employ high-frequency movements in OIS rates instead. Ampudia and Heuvel (2018) and Jarocinski and Karadi (2017) construct monetary policy shocks from movements in these derivatives.

The empirical results in this paper tie in with the results from two-agent models such as those in Galí, López-Salido, and Vallés (2007) and Bilbiie (2017), as well as richer

models by Werning (2015), Auclert (2019) and Hagedorn et al. (2019). A 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 and depends on the cyclical nature of income risk. If income risk is countercyclical², the models predict that a greater portion of constrained households increases the output response to a given monetary policy shock. 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 separately treat liquid and illiquid assets 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, Violante, and Weidner (2014), Kaplan and Violante (2014) and Kaplan, Moll, and Violante (2018).

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

II EFFECTS OF MONETARY POLICY SHOCKS ON OUTPUT

II.1 IDENTIFYING MONETARY POLICY SHOCKS

In order to estimate the effects of monetary policy to 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 2018a). More recently, Ampudia and Heuvel (2018) and Jarocinski and Karadi (2017) apply the approach to European data using Overnight Indexed Swap (OIS) rate movements around ECB announcements. These derivatives are traded over-

2. Guvenen, Ozkan, and Song (2014) show that income risk is countercyclical in the US.

the-counter between two parties exchanging a fixed interest rate for the floating Eonia overnight interest rate, both on a notional principal, for a pre-specified 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 (i) the fixed interest accrued on the principal and (ii) the interest accrued on the principal by investing it at the overnight interest rate every day is calculated and the contract is cash settled.³

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 (Jarocinski and Karadi 2017).⁴ 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 1999⁵. Figure I 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.

[Insert Figure I here]

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

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

4. We obtain time series on OIS rates at the minute frequency from Datastream. For more information, see Appendix A.

5. To 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).

fundamentals as the ECB. As pointed out by Jarocinski and Karadi (2017), 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 effect of the former and the latter.

We use the 3 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 weight 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

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

II.2 THE EFFECTS 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 (LP-IV) 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

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

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 regression⁶

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

As a benchmark we set the number of lags to $p = 3$ and the horizon of the impulse responses to $H = 36$.⁷ 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 .^{8,9} 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.¹⁰ 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 II presents impulse responses of real GDP for each country in our sample to an expansionary shock of 25 basis points. The IRFs are represented by the blue lines and, following Stock and Watson (2018), we construct 1 and 2 standard deviation confidence bands that surround the point estimates, using Newey-West estimators.

[Insert Figure II here]

The estimated impulse response functions reveal that expansionary monetary policy shocks cause output to increase in most countries. Output increases significantly after

6. For a detailed description of the data series used we refer the reader to Appendix C.

7. We 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.

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

9. Removing the set of control variables in (2), specially 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 24.99.

10. For 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.

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 1.8 percent after 27 months.

There is considerable heterogeneity in the responses, not only in magnitude, but also in the timing of the maximum impact. Moreover, the initial impacts of monetary policy shocks seem to be on average small and often not statistically different from zero.¹¹

The magnitudes of the responses are large, especially when compared to what is usually found in the literature when responses are estimated using standard VARs, i.e. without the use of external instruments (e.g., Christiano, Eichenbaum, and Evans 1999). This difference in magnitudes also arises when comparing single equation approaches to externally identified VARs, i.e. SVAR-IV. In Stock and Watson (2018), for example, the authors, using the LPIV approach, find that industrial production in the U.S. contracts by up to 9 percent following a contractionary shock of 100 basis points, while using a SVAR-IV approach the effects are significantly smaller, around 2 percent.¹²

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 countries that overlap with his sample, with a correlation coefficient of approximately 0.84.¹³ 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.

11. Although in general not statistically significant, the IRFs for some countries indicate a negative initial response to an expansionary monetary policy shock. One possible explanation is the existence of information shocks in addition to monetary policy shocks. In this case, the central bank's interest rate announcement also contains information about the future state of the economy. While it is beyond the scope of our paper, Jarocinski and Karadi (2017) and Nakamura and Steinsson (2018b) investigate this question. One exercise that we perform is to construct separate instruments based on market movements around the announcement at 13:45 and the press conference at 14:30, respectively. The main results of this paper hold when applying each of these. Under the assumption that information shocks are orthogonal to monetary policy shocks, using the instrument based on market movements around the announcement at 13:45 should give us unbiased IRFs. Leombroni et al. (2018) find that unexpected changes in the policy rate, which they call "target rate shocks", have significant effects on short-term rates, but little on long bond yields. We see this as suggestive evidence for shocks to the target rate being orthogonal to information shocks.

12. This difference in magnitudes and possible explanations are thoroughly explored in Coibion (2012).

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

III MEASURING FINANCIAL CONSTRAINTS

Bilbiie (2017) 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.

III.1 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. Two waves have been conducted, with data being released in 2013 and 2016. We rely on the second wave as the number of countries included is larger than in the first.

In approximating the share of households who have high MPC, we follow Kaplan, Violante, and Weidner (2014).¹⁴ A household is categorized as living 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 (Kaplan, Violante, and Weidner 2014).¹⁵

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 .¹⁶ We categorize a household as HtM if:

$$0 \leq m_i \leq \frac{y_i}{2} \tag{3}$$

14. 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.

15. U.S., Canada, Australia, U.K., Germany, France, Italy, Spain.

16. See Appendix B for details about the classification of assets as liquid and illiquid.

or if:

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

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 (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:

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

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.¹⁷

Figure IIIa 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).

[Insert Figure III here]

A concern with the measure described above is that households are interviewed during different points during the month or the 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 which 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

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

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 (categories (i) and (ii)) and label households that fulfill this criteria as being "Potentially Financially Vulnerable type 1" (PFV1).

The fractions are presented in Figure IIIb. Again, there is heterogeneity across countries and the average, indicated by the vertical line, is above 60%. 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 0.56, which we see as encouraging.

III.2 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 to measure 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 has covered all the countries in our sample, with a sample size of close to 90,000 households.¹⁸ 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 on 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.¹⁹ 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

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

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

and label it "Potentially Financially Vulnerable type 2" (PFV2). Figure IVa 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 0.68.

[Insert Figure IV here]

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.²⁰ 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 IVb. 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 HtM measure is 0.79.

IV LIQUIDITY CONSTRAINED HOUSEHOLDS AND MONETARY POLICY EFFECTIVENESS

IV.1 RESULTS

The results in section II.2 indicate that the countries in our sample do not respond homogeneously to monetary policy shocks. We proceed to link this finding with 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 three alternative measures introduced in Sections III.1 and III.2: PFV1, PFV2 and PFV3.

Scatterplots between different measures of monetary policy effectiveness and the shares of households living HtM are presented in Figure V. All panels show the share of HtM households on the horizontal axis and the vertical axes display different measures of the

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

effectiveness of monetary policy. Figure Va shows the peak of the output impulse response, which exhibits a significant positive correlation with the HtM share across countries. Figure Vb 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. Although the regression line in Figure Vc is subtly downward sloping, we do not view this as suggestive for a relationship between HtM and Time-to-Peak.

[Insert Figure V here]

We interpret these results in light of a standard TANK model as in Bilbiie (2017). 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 which occur due to the Euler equation of the unconstrained households. A shock which 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 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 V point to the latter, i.e. 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 (2017) if the income elasticity of constrained households with regards to aggregate income is larger than one. Furthermore, it is in line with the HANK model by Werning (2015) if income risk is countercyclical (a hypothesis for which Guvenen, Ozkan, and Song (2014) provide evidence).

Compared to Bilbiie (2017) and Werning (2015), Auclert (2019) allows for both varying maturities and different asset types. Although there are more channels through which different households can be differently exposed to changes in the nominal interest rate, he reaches a similar conclusion: aggregate output responses are larger if households with higher MPC are more affected.

We now turn to our alternative measures of MPC, namely PFV1-PFV3. Our focus is on peak responses, but as before, results are similar using cumulative responses as the

measure for monetary policy effectiveness. Time-to-peak is not significantly correlated with the measures that we explore below (not shown).

The three scatterplots are presented in Figure VI. Correlations between the peak responses and each of the three statistics PFV1-PFV3 are strong. We view this as encouraging for two reasons. First, the results lend credence to the measure proposed by Kaplan, 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 (b) and (c) 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.

[Insert Figure VI here]

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. With this in mind, Figure VII shows the relationship between the peak responses across countries and their share of wealthy (panel (a)) and poor (panel (b)) HtM households. 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. Panel (c) in Figure VII gives no evidence for a relationship between output responses and

the share of asset poor and the statistic is outperformed by all of our other measures in predicting by how much output is affected through monetary policy shocks.

[Insert Figure VII here]

IV.2 ROBUSTNESS

First, we show that our results are robust to changing the horizon at which the effects of monetary policy are measured. In the previous section, 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 then correlate each of these with the HtM values.

[Insert Figure VIII here]

The horizontal axis in Figure VIIIa measures the horizon 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 towards the end of our estimation horizon. The latter, again, is unsurprising, as Figure II 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.²¹

Figure VIIIb graphs the results from when we estimate 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 is a relevant statistic for the effectiveness of monetary policy across countries.

Next, we investigate whether other country-specific aggregates can account for the heterogeneity in impulse responses that we observe. We focus on a set variables that could

21. 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.

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 cross country perspective. Subsequently, for each variable we investigate whether (i) it is correlated with HtM shares, (ii) it is correlated with the peak effects we find in section II and (iii) whether after controlling for the variable, the HtM share still explains part of the output responses we observe.²²

Cloyne, Ferreira, and Surico (2018) find that households who 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 house prices and output via the collateral channel (see Cloyne et al. 2017). 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 residence but do not have a mortgage attached to it. We label the variable for the fraction of these households in each country as *HO*.

It is possible that the the effectiveness of monetary policy depends on how highly indebted households are (e.g., Flodén et al. 2017) and on how common it is that mortgages have an adjustable interest rate (e.g., Calza, Monacelli, and Stracca 2013; Flodén et al. 2017). We calculate the fraction of households that have *at least one* mortgage with an adjustable interest rate and label the variable *Flex*. To test if 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.

22. 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.

In Section 5.1. we argued that it is mainly the fraction of wealthy HtM households that explains why the total fraction of HtM households is correlated with peak values. A possibility is 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 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*.²³

Wong (2019) finds that, in the U.S., 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 households, on average across the countries, are less likely to be HtM when age increases. Moreover, the probability of being wealthy HtM increases between age 20 and the late 30's, and decreases after this threshold. We find it important to test if 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 Global VAR (GVAR) models (e.g., Burriel and Galesi 2018; Georgiadis 2015) emphasizes the importance of considering spillover effects of monetary policy and the size of these spillovers are 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 if 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 respond 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 which are half as large. We have mentioned that our estimates for the effectiveness of monetary policy

23. A household whose net illiquid assets are positive is labelled as wealthy.

are similar to the estimates in his paper. Hence 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 years 2000-2012 for each country.

All results are summarized in Table I. The first column in the table 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 magnitudes with both peak effects and HtM shares. Another is *Services* which is negatively correlated with peak effects (confirming the result from Georgiadis 2015) and also negatively correlated with HtM shares.

[Insert Table I here]

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 for why we find such a strong correlation between peak responses and HtM shares. To get a sense for 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 in Table I and tell us the correlation coefficient between the peak effects and HtM shares, after that the variation in HtM shares explained by these other variables, varying across the rows in the table, has been taken out. 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 extract 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 I, we find nothing that supports that the correlation between the peak effects and HtM shares is caused by omitted variables.

V 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 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), since standard VAR approaches are infeasible for large groups of countries. 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, peak values and timing.

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 three 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 empirical work, given plausible assumptions about income risk (Werning 2015) or the elasticity of constrained households' incomes to aggregate income (Bilbiie 2017).

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.

Mattias Almgren, Jose E. Gallegos, John Kramer and Ricardo Lima are affiliated with the IIES, Stockholm University.

References

- Ampudia, Miguel, and Skander Van den Heuvel. 2018. “Monetary policy and bank equity values in a time of low interest rates.” *ECB Working Paper*.
- Auclert, Adrien. 2019. “Monetary policy and the Redistribution Channel.” *American Economic Review*, forthcoming.
- Bilbiie, Florin. 2017. “The New Keynesian Cross: Understanding Monetary Policy with Two Agents.” *mimeo*.
- Botero, Juan C, Simeon Djankov, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer. 2004. “The regulation of labor.” *The Quarterly Journal of Economics* 119 (4): 1339–1382.
- Burriel, Pablo, and Alessandro Galesi. 2018. “Uncovering the heterogeneous effects of ECB unconventional monetary policies across euro area countries.” *European Economic Review* 101:210–229. ISSN: 00142921.
- Calza, Alessandro, Tommaso Monacelli, and Livio Stracca. 2013. “Housing finance and monetary policy.” *Journal of the European Economic Association* 11 (suppl_1): 101–122.
- 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, C., and A. Lin. 1971. “Best Linear Unbiased Interpolation , Distribution , and Extrapolation of Time Series by Related Series.” *The Review of Economics and Statistics* 53 (4): 372–375.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1999. “Chapter 2 Monetary policy shocks: What have we learned and to what end?” *Handbook of Macroeconomics* 1 (PART A): 65–148. ISSN: 15740048.
- Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2018. “Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism.” *Mimeo*.

- Cloyne, James, Kilian Huber, Ethan Ilzatzki, and Henrik Kleven. 2017. “The effect of house prices on household saving : A new approach.” *NBER working paper series*.
- Coibion, Olivier. 2012. “Are the effects of monetary policy shocks big or small?” *American Economic Journal: Macroeconomics* 4 (2): 1–32.
- 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.” *Mimeo*.
- Flodén, Martin, Matilda Kilström, Josef Sigurdsson, and Roine Vestman. 2017. “Household Debt and Monetary Policy : Revealing the Cash-Flow Channel.” *Mimeo*.
- 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): 227–270.
- Galí, Jordi, and Tommaso Monacelli. 2008. “Optimal monetary and fiscal policy in a currency union.” *Journal of International Economics* 76 (1): 116–132. ISSN: 00221996.
- 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 (January): 195–215. ISSN: 00142921.
- Gerlach, Stefan, and Frank Smets. 1995. “The Monetary Transmission Mechanism: Evidence from the G-7 Countries.” *BIS Working Papers* 26.
- Gertler, Mark, and Peter Karadi. 2015. “Monetary Policy Surprises, Credit Costs, and Economic Activity.” *American Economic Journal: Macroeconomics* 7 (1): 44–76. ISSN: 1945-7707.
- Güvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. “The Nature of Countercyclical Income Risk.” *Journal of Political Economy* 122 (3): 621–660.
- Hagedorn, Marcus, Jinfeng Luo, Iourii Manovskii, and Kurt Mitman. 2019. “Forward guidance.” *Journal of Monetary Economics*.

- Honkkila, Juha, and Ilja Kristian Kavonius. 2013. “Micro and Macro Analysis on Household Income, Wealth and Saving in the Euro Area.” *ECB Working Paper Series* 1619 (1619).
- Jarocinski, Marek, and Peter Karadi. 2017. “Central Bank Information Shocks.” *Mimeo*.
- Jordà, Òscar. 2005. “Estimation and Inference of Impulse Responses by Local Projections.” *American Economic Review* 95 (1): 161–182. ISSN: 00028282.
- Kaplan, Greg, Kurt Mitman, and Giovanni Violante. 2017. “The Housing Boom and Bust: Model Meets Evidence.” *Mimeo*. ISSN: 0007-666X.
- 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* 2014 (1): 77–138. ISSN: 1533-4465.
- Kuttner, Kenneth N. 2001. “Monetary policy surprises and interest rates: Evidence from the Fed funds futures market.” *Journal of Monetary Economics* 47 (3): 523–544. ISSN: 03043932.
- Leombroni, Matteo, Andrea Vedolin, Gyuri Venter, and Paul Whelan. 2018. “Central bank communication and the yield curve.” *Available at SSRN 2873091*.
- Lloyd, Simon. 2018. “Estimating Nominal Interest Rate Expectations: Overnight Indexed Swaps and the Term Structure.” *Bank of England Working Paper*.
- Nakamura, Emi, and Jón Steinsson. 2018a. “High-frequency Identification of Monetary Non-neutrality: The Information Effect.” *The Quarterly Journal of Economics* 133 (3): 1283–1330.
- . 2018b. “Identification in macroeconomics.” *Journal of Economic Perspectives* 32 (3): 59–86.
- Ramey, V. A. 2016. “Macroeconomic Shocks and Their Propagation.” *Handbook of Macroeconomics* 2:71–162. ISSN: 15740048.

- Stock, James H, and Mark W Watson. 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *The Economic Journal* 128 (610): 917–948.
- Werning, Iván. 2015. "Incomplete markets and aggregate demand." *Mimeo*.
- Wong, Arlene. 2019. "Refinancing and the Transmission of Monetary Policy to Consumption." *Mimeo*.

APPENDIX

A EUROPEAN OVERNIGHT INDEXED SWAP DATA

We obtain a minute-frequency series for Eonia Overnight Indexed Swaps from Datasream. We compute the fixed rate of the swap as the mid point between the bid and ask price at the close of each minute. We then drop all dates from the sample that are not ECB announcement dates.

The resulting series contains implausible outliers, e.g. the rate decreasing to zero for one minute, or short fluctuations of more than 5 standard deviations. Consequently, we drop the highest and lowest percentile of observations on each announcement day. Lastly, we manually drop remaining implausible observations if they fall within either of the two announcement windows.

B CALCULATE HTM SHARES USING DATA FROM THE HFCS

The HFCS imputes data for missing values related to assets, liabilities and income variables. Our calculations are partly based on these imputed data. A missing value is imputed five times (multiple imputation), where each time a different random term is added to the predicted value. If this would not be done, imputation uncertainty would not be taken into account. This has the consequence that statistics can vary between implicates (the differences are in general very small). Whenever we calculate statistics based on HFCS data, we average over all the implicates. We consistently use the cross-sectional weights, which are mainly intended to compensate for some households being more likely to be selected into the sample than others.

Before we label households, we drop observations where the age of the reference person in the household is below 20 or above 79. As in Kaplan, Violante, and Weidner (2014) we drop observations when the only income that the household receives is from self-employment. The results do not change markedly if we choose to keep these observations.

We need to categorize variables as liquid wealth, illiquid wealth, liquid debt and illiq-

uid debt. We follow Kaplan, Violante, and Weidner (2014) to a large extent. In the tables below we present what variables go into respective category and the *Name* refers to its unique name in the HFCS data. The difference between how we categorize the variables and how Kaplan, Violante, and Weidner (2014) do it is that we categorize saving accounts as liquid assets while they categorize it as illiquid for the European countries. We choose to categorize it as liquid as it is our view that households can, in general, make adjustments to the balance on saving accounts without incurring substantial costs. In the Panel Study of Income Dynamics (PSID), saving accounts are combined with other assets such as checking accounts. Moreover, in the calibration of the models in Kaplan, Mitman, and Violante (2017) and Carroll et al. (2017), saving accounts are categorized as liquid.

In the calculation of HtM shares, Kaplan, Violante, and Weidner (2014) assume that households on average are paid bi-weekly. In our calculations we will assume that households on average are paid once every month, which we believe is a more accurate assumption about the payment frequency in European countries. We define liquid wealth = liquid assets-liquid debt and illiquid wealth = illiquid assets-illiquid debt.

[Insert Table A.1 here]

C LOCAL PROJECTION DATA

Inflation: We obtain the monthly Harmonized Index of Consumer Prices for all items for all countries in our sample and the euro area from Eurostat (*prc_hicp_midx*).

Industrial Production: We obtain monthly values for Industrial Production (excluding construction) from Eurostat. The series is seasonally and calendar adjusted (*sts_inpr_m*). Because Ireland changed its formula for the calculation of some national aggregates, we make some assumptions to keep the series as coherent as possible. The change affects the value of Industrial Production in the first two months of 2015, resulting in growth rates in excess of 10%. We substitute these two growth rates with the average growth over 2014, which results in a level shift for all IP values after March 2015.

Unemployment rate: We obtain monthly values of the unemployment rates for all countries in our sample from Eurostat (*une_rt_m*). The rates are measured for the active pop-

ulation aged 25 to 74 and are seasonally and calendar adjusted. For Estonia, the value of January 2000 is missing. We obtain it from the OECD (*LRHUADTT*). The rest of the series coincides with the values from Eurostat.

Real GDP: We obtain the quarterly values for Real GDP for all countries in our sample from Eurostat (*namq_10_gdp*). The series measures chain-linked volumes of Gross Domestic Product and is seasonally and calendar adjusted. Again, we adjust the series of Ireland due to implausibly high GDP growth in the first quarter of 2015. We substitute the reported growth rate in 2015Q1 with the average growth rate during 2014, which results in a level shift of all subsequent observations.

Eonia: We obtain values for the European OverNight Index Average from Eurostat (*irt_st_m*).

Composite Index of Systematic Stress: We obtain values for the SCISS for all countries in our sample and the euro area as a whole from the ECB . The data frequency is monthly (*CISS.M..Z0Z.4F.EC.SOV_CI.IDX*).

Retail trade: We obtain monthly data on Retail trade, except of motor vehicles and motorcycles from Eurostat for all countries in our sample. The series refers to deflated turnover and is seasonally and calendar adjusted (*sts_trtu_m*).

Bond yields: We obtain EMU convergence criterion bond yields for all countries in our sample from Eurostat (*irt_lt_mcby_m*).

D TABLES

Table I: Correlations and semipartial correlations

X	$\rho(\text{Peak}, X)$	$\rho(\text{HtM}, X)$	$\rho(\text{Peak}, \text{HtM} - X)$
<i>Peak</i>	1.00 (0.000)	0.78 (0.000)	NA (NA)
<i>HtM</i>	0.78 (0.000)	1.00 (0.000)	NA (NA)
<i>Own</i>	0.42 (0.086)	0.36 (0.143)	0.68 (0.003)
<i>Mort</i>	-0.35 (0.156)	-0.32 (0.198)	0.71 (0.001)
<i>HO</i>	0.54 (0.022)	0.47 (0.047)	0.60 (0.011)
<i>Wealthy</i>	0.35 (0.149)	0.24 (0.346)	0.72 (0.001)
<i>Flex</i>	-0.04 (0.893)	-0.03 (0.924)	0.79 (0.000)
<i>Age</i>	-0.05 (0.840)	0.09 (0.726)	0.79 (0.000)
<i>LTV</i>	-0.34 (0.165)	-0.26 (0.304)	0.72 (0.001)
<i>LTI</i>	-0.36 (0.146)	-0.22 (0.375)	0.72 (0.001)
<i>Trade</i>	0.10 (0.67)	-0.18 (0.474)	0.81 (0.000)
<i>ROL</i>	-0.08 (0.797)	0.11 (0.716)	0.88 (0.000)
<i>Services</i>	-0.41 (0.083)	-0.20 (0.438)	0.73 (0.001)

Note: 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 Households Finance and Consumption Survey (HFCS). See text for information about the source of the other variables.

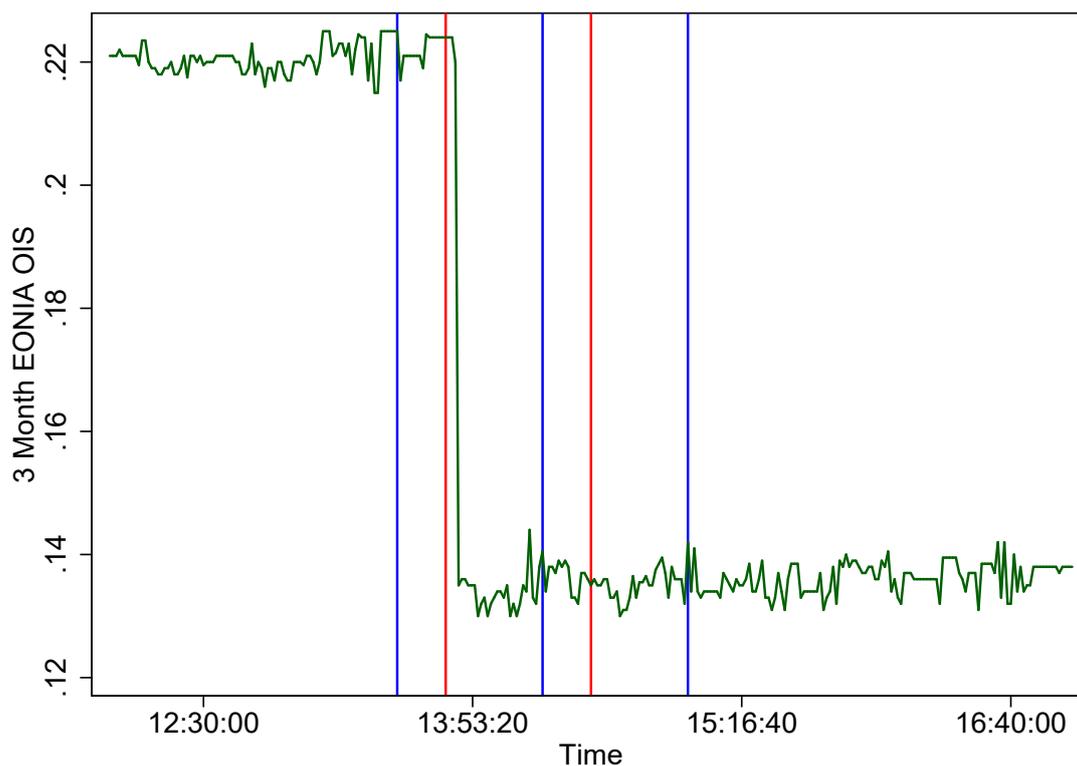
Table A.1: Classification of Income and assets in the HFCS

Name	Description	comment
Income		
di1100	employee income	
di1610	unemployment benefits	
di1620	other social transfers	
hg0210	income from regular private transfers	
di1510	gross income from public pensions	
Liquid assets		
hd1110	value of sight accounts	
da2102	mutual funds, total	
da2105	shares, publicly traded	
da2103	bonds	
hd1210	value of saving accounts	
Illiquid assets		
hb0900	current price of household main residence	
hb280x	other property \$x: current value	
hb2900	additional properties current value	
sum of pf0710 across HH members	current value of all occupational pension plans that have an account	
da2109	voluntary pension/whole life insurance	
Liquid debt		
hc0220	amount of outstanding credit line/overdraft balance	
hc0320	amount of outstanding credit cards balance	
Illiquid debt		
hb170x	HMR mortgage \$x: amount still owed	x = {1,2,3}
hb370x	other property mortgage \$x: amount still owed	x = {1,2,3}
hb4100	money still owed on additional other property loans	
hb2100	money still owned on additional HMR loans	
hc080x	non-collateralised loan \$x: outstanding balance of loan	x = {1,2,3}
	hc080x counted if $hc050xa \leq 2$	

a. This adjustment follows Kaplan, Violante, and Weidner 2014.

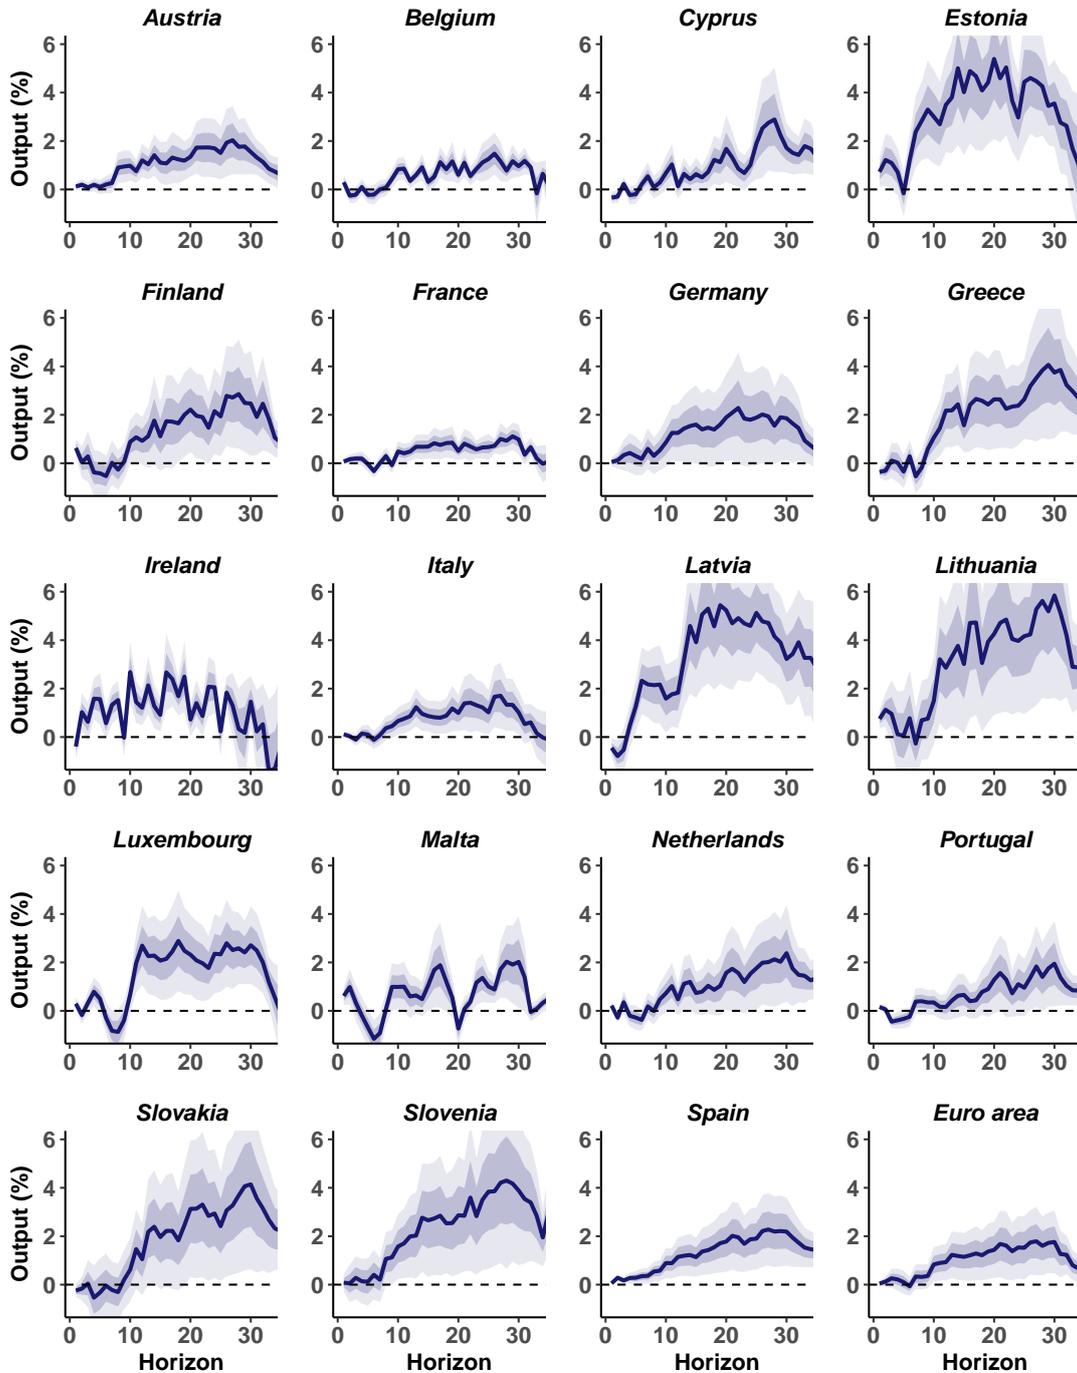
E FIGURES

Figure I: Overnight Indexed Swap rates on 05.07.2012



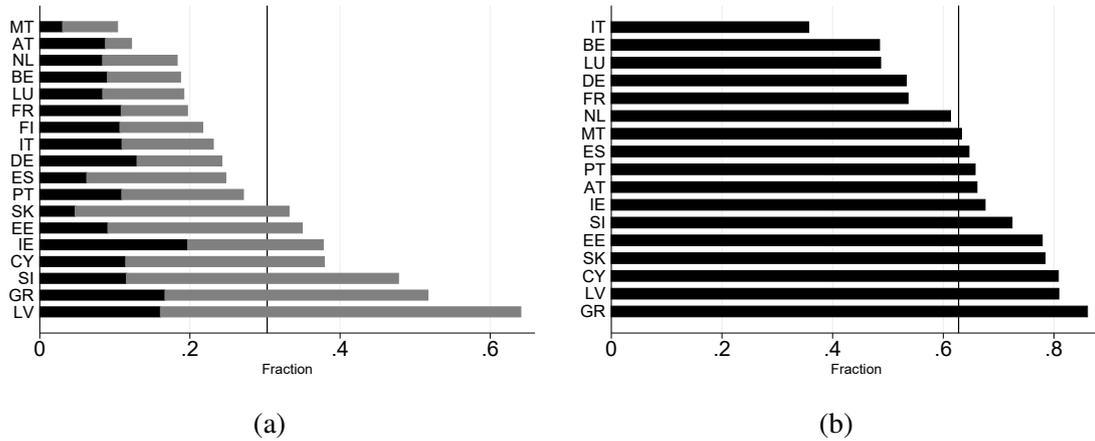
Note: This figure shows the time series of the 3 month EONIA Overnight Indexed Swap for July 5, 2012. The blue lines represent the borders of our measurement windows, the red lines indicate the policy events, i.e. the ECB's press release at 13:45 CET and the start of the press conference at 14:30 CET. The first pre-window runs from 13:30-13:44 CET and then the first post-window is active between 13:45-14:14 CET. Then a second pre-window runs from 14:15-14:29 CET and the second post-window is active between 14:30-14:59 CET.

Figure II: Impulse responses for output in euro zone countries – LPIV



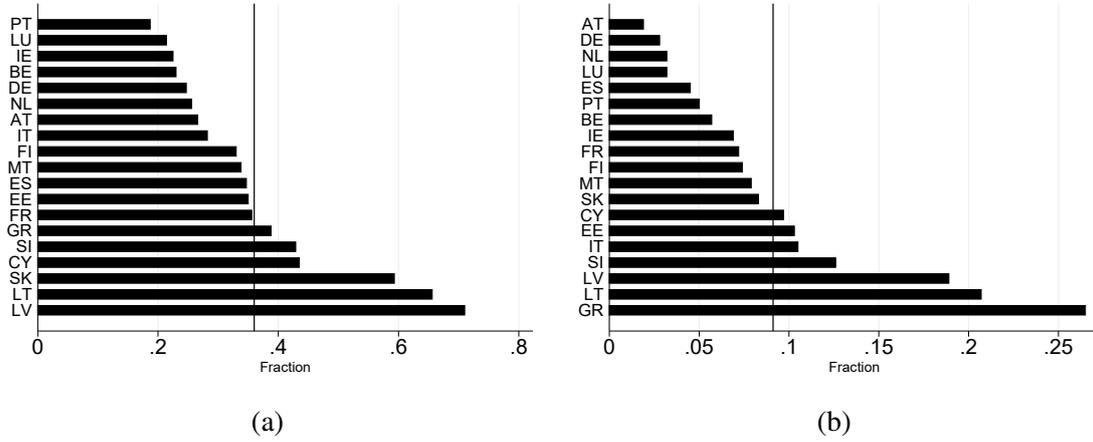
Note: This figure shows impulse responses of real GDP to a 25 basis points expansionary monetary policy shock. For each euro area country, the response is estimated using LPIV (Equation 1). The solid blue lines represent the IRFs produced by our preferred specification (see text for details). The dark and light blue shaded areas represent 1 and 2 standard deviation confidence bands, constructed using Newey-West estimators.

Figure III: Wealthy HtM, Poor HtM and PFV1



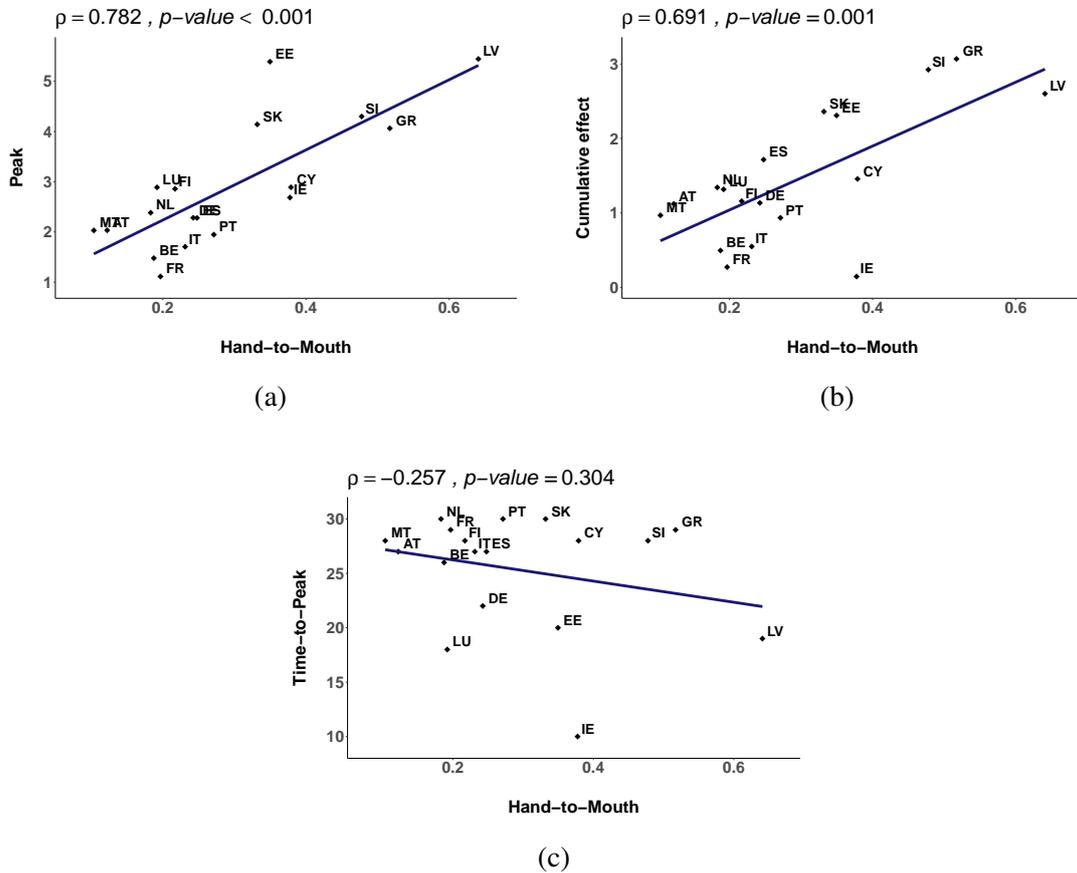
Note: *Panel (a):* This figure shows the fraction of households that are classified as HtM in each country. The total fraction, given by the total length of each bar, is divided up into two parts: poor (black) and wealthy (gray). The vertical line indicates the unweighted average of total HtM in our sample of countries. We do not have data for Lithuania. We use data from the Households Finance and Consumption Survey (HFCS). *Panel (b):* The figure 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 is missing for Finland and Lithuania, hence they are not included in the figure. We use data from the Household Finance and Consumption Survey (HFCS).

Figure IV: PFV2 and PFV3



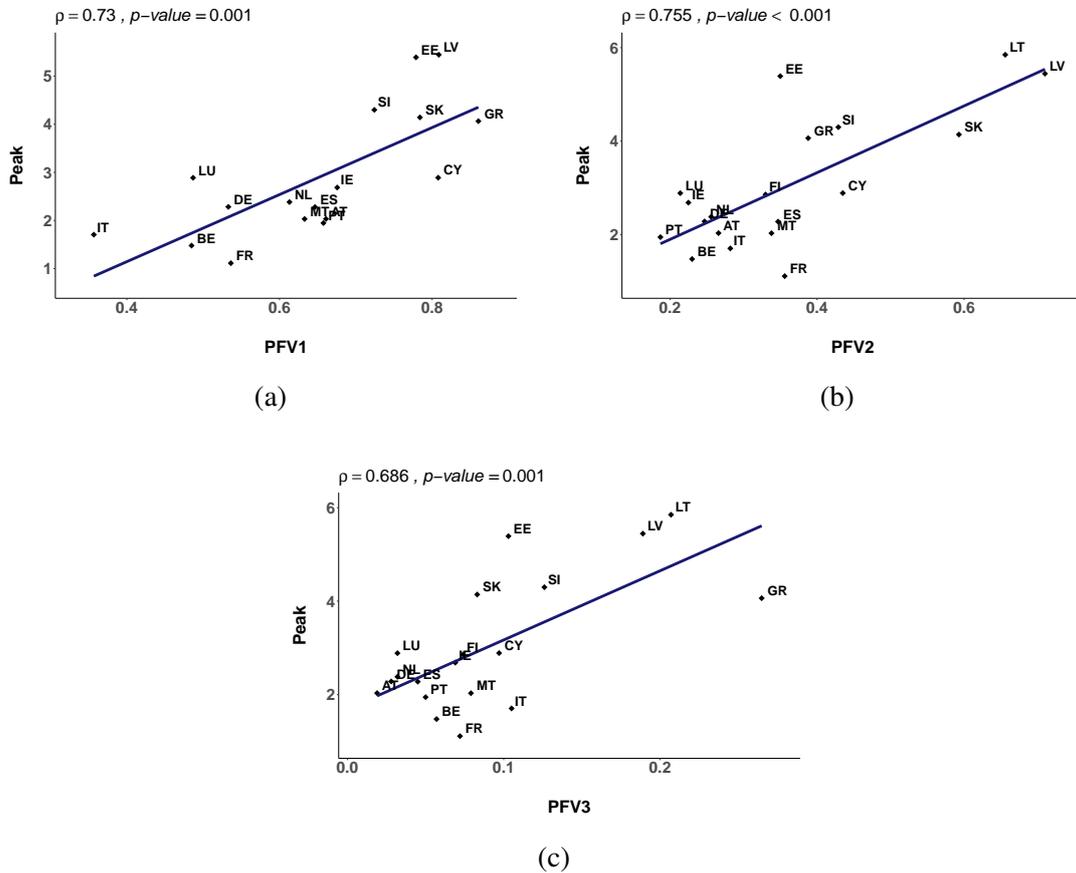
Note: *Panel (a):* The figure shows the fraction of households that believe that they are unable to face unexpected expenses with the use of own resources (PFV2). The fraction is given by the length of each bar. The vertical line indicates the average of the fractions in our sample of countries. Data is from European Union Statistics on Income and Living Conditions (EU-SILC). *Panel (b):* The figure 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 is from European Union Statistics on Income and Living Conditions (EU-SILC).

Figure V: Monetary policy effectiveness and Hand-to-Mouth shares



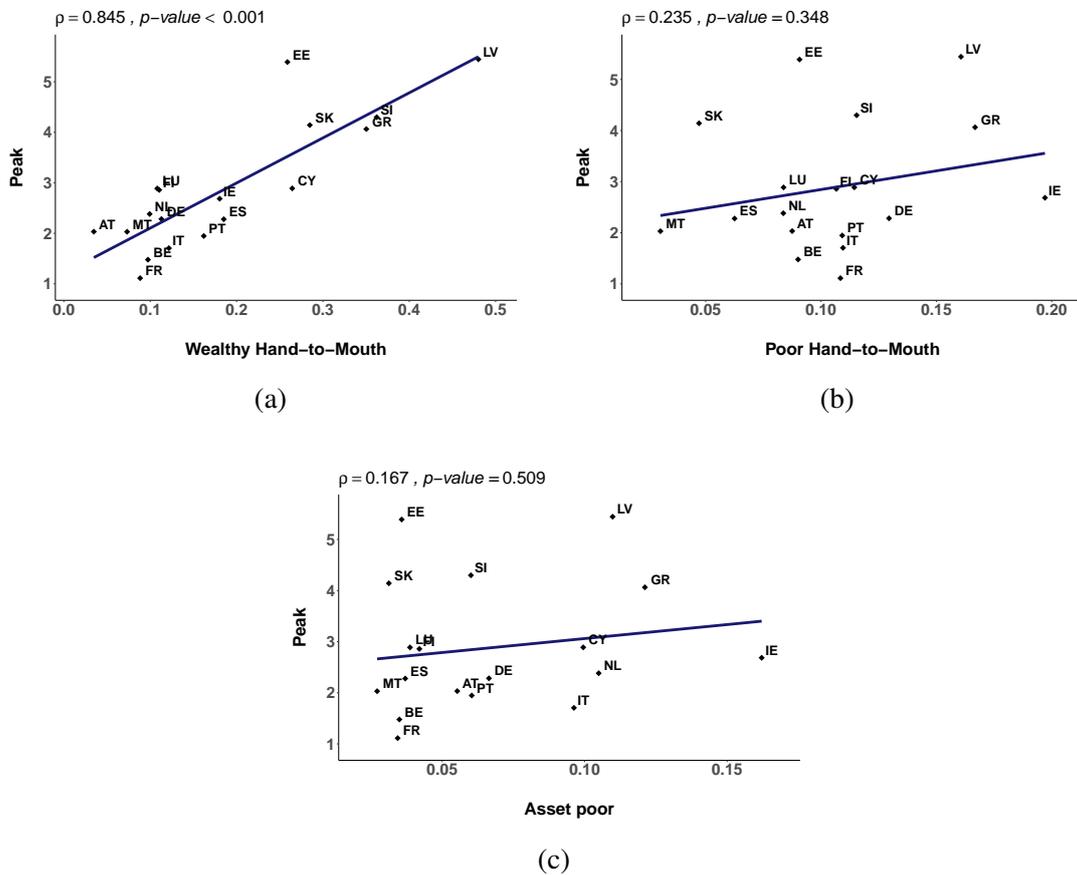
Note: This figure plots the effectiveness of monetary policy, as measured by the peak effect, cumulative effect and time-to-peak from the real GDP impulse responses, calculated using the benchmark LPIV estimation, against the share of households classified as living HtM in each euro country (except Lithuania, not included in the HFCS). The HtM shares are calculated using data from the Households Finance and Consumption Survey (HFCS). The impulse is an expansionary 25 basis point monetary policy shock. The blue lines are fitted from regressions of Peak/Cumulative/Time-to-Peak values on HtM shares. In the upper left corner of each panel we report the correlation coefficient ρ and the p-value. *Panel (a)*: Peak effects and share of Hand-to-Mouth. *Panel (b)*: Cumulative effects and share of HtM. *Panel (c)*: Time-to-peak and share of HtM.

Figure VI: Impact of monetary policy and alternative liquidity constraint measures



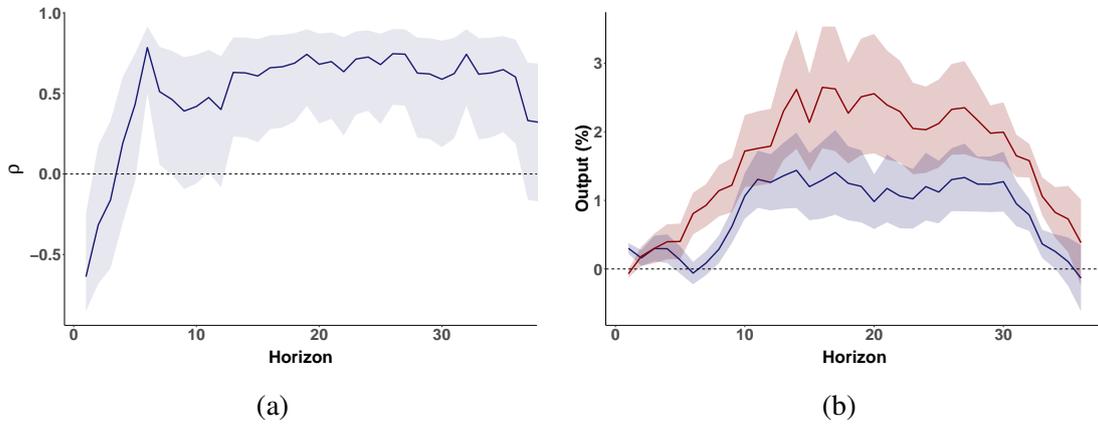
Note: This figure plots the effectiveness of monetary policy, as measured by the peak effect, calculated using the benchmark LPIV estimation, against different statistics in each euro area country (except Lithuania in the bottom-right panel). The impulse is an expansionary 25 basis point monetary policy shock. The blue lines are fitted from a regression 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):* Peak effects and PFV1. *Panel (b):* Peak effects and PFV2. *Panel (c):* Peak effects and PFV3. PFV1-PFV3 as defined in Section III. PFV1 is calculated using data from the Households Finance and Consumption Survey (HFCS) and to calculate PFV2 and PFV3 we use data from European Union Statistics on Income and Living Conditions (EU-SILC).

Figure VII: Monetary policy effectiveness and Wealthy and Poor Hand-to-Mouth shares



Note: 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 (Kaplan, Violante, and Weidner 2014), respectively, in each euro area country (except Lithuania). The impulse is an expansionary 25 basis point monetary policy shock. The blue lines are fitted from regressions of peak values on Wealthy or Poor HtM shares. In the upper left corner of each panel we report the correlation coefficient ρ and the p-value. *Panel (a):* Peak effects and share of Wealthy Hand-to-Mouth. *Panel (b):* Peak effects and share of Poor HtM. *Panel (c):* Peak effects and asset poor. Wealthy HtM shares, Poor HtM shares and shares of assets poor are calculated using data from the Households Finance and Consumption Survey (HFCS).

Figure VIII: Monetary policy effectiveness and Wealthy and Poor Hand-to-Mouth shares



Note: *Panel (a)*: correlation between HtM and peak responses, retrieved from the impulse responses using the benchmark specification, for each horizon $h = \{0, 1, \dots, H\}$. The shaded area is the 95% confidence band around the point estimates. *Panel (b)*: 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% confidence bands around the point estimates. Calculations of HtM shares are based on data from the Households Finance and Consumption Survey (HFCS).