How Does Monetary Policy Affect Income and Wealth Inequality? Evidence from Quantitative Easing in the Euro Area

February 14, 2019

Michele Lenza[1]
European Central Bank and ECARES–ULB

Jiri Slacalek[2]
European Central Bank

Abstract

This paper evaluates the impact of quantitative easing on income and wealth of individual euro area households. We first estimate the aggregate effects in a VAR model with unemployment, wages, interest rates, house prices and stock prices. We then distribute the aggregate effects across households using a reduced-form simulation on micro data, which captures the portfolio composition, the income composition and the earnings heterogeneity channels of transmission. The earnings heterogeneity channel is important: QE compresses the income distribution since many households with lower incomes become employed. In contrast, monetary policy has only negligible effects on the Gini coefficient for wealth.

Keywords Monetary Policy, Household Heterogeneity, Inequality, Income, Wealth, Quantitative Easing, Great Recession

JEL codes D14, D31, E44, E52, E58

[1]Lenza: DG Research, European Central Bank and ECARES–ULB, email: michele.lenza@ecb.europa.eu
[2]Slacalek: DG Research, European Central Bank, http://www.slacalek.com/, email: jiri.slacalek@ecb.europa.eu

This paper uses data from the Household Finance and Consumption Survey. We thank Marco Felici and Thibault Cézanne for excellent research assistance and Lídia Brun, Vítor Constâncio, Maarten Dossche, Michael Ehrmann, Nicola Fuchs-Schündeln, Jordi Galí, Bertrand Garbinti, Dimitris Georgarakos, Domenico Giannone, Marek Jarociński, Bartosz Maćkowiak, Alberto Martin, Klaus Masuch, Pierre Monnin, Giorgio Primiceri, Moreno Roma, Alfonso Rosolia, Anna Samarina, Oreste Tristani, Mika Tujula, Panagiota Tzamourani, Gianluca Violante, Bernhard Winkler, Christian Wolf and Jonathan Wright, and seminar audiences at ECB, HFCN, Barcelona GSE Summer Forum and 2018 IAAE Annual Conference for useful comments and discussions. The views presented in this paper are those of the authors, and do not necessarily reflect those of the European Central Bank.

1 Introduction

The collection of reliable data in recent years has allowed researchers to characterize the evolution of wealth and income distributions in time and across countries. In particular, Piketty (2013) shows that, contrary to the traditional view based on Kuznets (1955), advanced economies do not inevitably evolve toward more egalitarian societies. This fact has sparked an intense debate about the drivers of economic inequality. In general, inequality is seen as related to the dynamics of the structural features of economies, such as the emergence of skill-biased technological progress (Katz and Murphy, 1992; Acemoglu, 2002; Autor, 2014), the deepening of globalization (Katz and Autor, 1999), the tendency toward the reduction in the progressivity of tax systems (Alvaredo et al., 2013) and portfolio heterogeneity (Fagereng et al., 2016; Hubmer et al., 2018).

Recently, since central banks have started to undertake extensive asset purchase programmes to circumvent the lower bound on nominal interest rates, monetary policy has also been put forth as a possible driver of economic inequality.[1] This paper investigates how unconventional monetary policy, specifically the quantitative easing (QE) program of the European Central Bank,[2] affects the distribution of income and wealth across individual households in the euro area. The analysis proceeds in two steps, making use of both aggregate and household-level data. First, we use aggregate data to estimate the effects of QE on unemployment, income and asset prices. Second, we distribute the aggregate effects across individual households using the information on their assets and income.

In more detail, in the first stage we estimate the transmission mechanism of a euro area QE shock. Since monetary transmission may differ across countries, we specify a large multi-country VAR model which includes both euro area and country-specific variables from the four largest euro area countries (France, Germany, Italy and Spain). The euro area variables cover most notably short-term and long-term interest rates, on which our strategy to identify monetary policy shocks partly hinges. The main identifying assumption for the QE shock is that it generates a negative correlation between the term spread (defined as long-term minus short-term interest rate) and real GDP in the four countries. The country-specific variables include, among others, those related to the dynamics of household income and wealth: the unemployment rate, wages and house prices.[3] Allowing for cross-country heterogeneity in the transmission mechanism is important, as the impulse responses of unemployment rates and asset prices vary across

[1] See Colciago et al. (2018) for a comprehensive survey of the theoretical and empirical literature on the effects of conventional and unconventional monetary policy on inequality.

[2] By quantitative easing of the ECB we mean the Asset Purchase Programme (APP), which started in January 2015 in order to address the risks of a long period of low inflation. The APP includes various purchase programmes under which private sector securities and public sector securities (including sovereign bonds) are bought. For an early assessment of the APP see Andrade et al. (2016).

[3] The large dimension of the model (25 variables) and the relatively short available sample (1999Q1 to 2016Q4) is handled using Bayesian estimation methods with informative priors which, as suggested by De Mol et al. (2008) and Bańbura et al. (2010), controls for overfitting while at the same time extracting the valuable information in the sample. The informativeness of the prior distributions is set according to the hierarchical BVAR procedure developed in Giannone et al. (2015). For the identification of the QE shocks, we impose a combination of zero and sign restrictions exploiting the method described in Arias et al. (2018), borrowing some elements of the identification scheme in Baumeister and Benati (2013).

countries: for example, the unemployment rate in Spain responds considerably more to the QE shock than in Germany. The comparison with alternative estimates also shows that existing studies find similar aggregate effects for the real and nominal variables in our model concluding, in general, that asset purchase programs such as QE have noticeable effects on the real economy (for an extensive recent survey, see Dell'Ariccia et al., 2018, and Tables C.1 and C.2 in the Online Appendix C).

However, aggregate cross-country heterogeneity is not the only relevant dimension to capture the different impact of QE across households. Indeed, the aggregate effects may result in heterogeneous impacts on households also because of the substantial differences in their sources of income (e.g., employment status, labor vs financial income) and their portfolios (holdings of real estate, shares and bonds). Consequently, in the second stage, we distribute the aggregate effects estimated in the VAR across the individual households using micro data on the composition of their assets and income. The analysis in this part relies on the Household Finance and Consumption Survey (HFCS), a dataset which collects detailed household-level information on balance sheets, income and socio-demographic variables for European countries in a similar way the Survey of Consumer Finances does for the US. Our analysis captures the transmission of QE to households via three channels: (i) income composition, (ii) portfolio composition and (iii) earnings heterogeneity.

The two composition channels operate via the heterogenous reaction of various income and wealth components to monetary policy. Figure 1 shows that the share of key income components varies substantially with the level of household income. Households in the lowest income quintile earn only roughly 20 percent of their gross income as employee income, while those in the top quintile about 60 percent. Similarly, the share of financial and rental income increases from 2 percent to almost 10 percent. In contrast, the share of transfers and unemployment benefits declines across income quintiles from almost 20 percent to about 3 percent. Figure 2 documents that the composition of household wealth is similarly varied. For example, the share of self-employment business wealth and stock market wealth (shares) on total assets in the top net wealth quintile is substantially larger, while the share of real estate is lower. To empirically capture the two composition channels, we update the components of income and wealth at the household level using the aggregate impulse responses for wages and for house, stock and bond prices.[4]

The earnings heterogeneity channel, instead, consists of the heterogeneous reaction of the employment status and hours worked to monetary policy. To capture this channel, we run a reduced-form simulation (as in Ampudia et al., 2016) which redistributes the aggregate decline in unemployment across individuals depending on their demographic characteristics: some unemployed individuals become employed and receive a substantial increase in (labor) income, as they start earning wages rather than unemployment benefits. The simulation ensures that the reduction of the unemployment rate in the household data is consistent with the aggregate drop in unemployment in the VAR impulse responses.

[4] In the baseline setup we assume that household portfolios are not rebalanced in response to the announcement of QE. This assumption is supported by the empirical evidence on considerable inertia in household portfolios, e.g., Ameriks and Zeldes (2004), Brunnermeier and Nagel (2008), Andersen et al. (2018) and others.

Our empirical results show that accounting for household heterogeneity in income and wealth is indeed important for describing the effects of quantitative easing on income and wealth inequality. For income, the overall effect of quantitative easing is dominated by the earnings heterogeneity channel: transitions from unemployment to employment account for about 75% of the effect on mean income across households. Importantly, the contribution of this channel is particularly pronounced in the lower part of the income distribution. One year after the occurrence of an exogenous QE-shock driving down on impact the term spread by 30 basis points, the unemployment rate among households in the bottom income quintile declines by 2 percentage points and mean income increases by more than 3 percent, accounting for more than 90% of the total increase in income.[5]

Overall, QE reduces income inequality via the earnings heterogeneity channel, while the income composition channel increases more incomes at the top, but is substantially smaller. Summing the effects of the two channels just described, QE noticeably compresses the income distribution: the Gini coefficient for gross household income declines from 43.1 to 42.9 percent, one year after the shock. While the effects are likely to fade away over longer horizons, given the likely transient nature of the effects of monetary policy, this evidence suggests that quantitative easing contributes to support vulnerable households, mainly via the earnings heterogeneity channel. Our main robustness checks pertain to alternative scenarios in which financial income strongly increases due to QE. While the increase in financial income is particularly beneficial for the top tail of the income distribution, its contribution to the changes in total income is limited and it does not significantly change our results on income inequality.

We then investigate how QE changes the wealth distribution via the portfolio composition channel. The policy temporarily increases the value of stocks and self-employment businesses, both mostly held by wealthier households. However, our estimates of the effects on net wealth are essentially driven by housing wealth, which reflects the fact that 60% of euro area households own their main residence and, overall, real assets account for about 70–80 percent of total assets across the wealth distribution. As expected, the effects of quantitative easing on net wealth tend to be stronger for leveraged households, relatively to their wealth level. At the same time, poorer households have a lower level of wealth and the effects of QE relative to the wealth level do not immediately translate in the effects on inequality. To gauge the latter, once again we compute the change in the Gini index implied by the effects of QE on asset prices and find that inequality in the net wealth distribution declines, but only by a negligible amount. This conclusion remains unaffected if we allow for some rebalancing of financial portfolios and for more differentiated responses of house prices to QE.

Our paper is related to the growing literature on the effects of monetary policy on inequality. Coibion et al. (2017) use quarterly data from the US Consumer Expenditure Survey in a VAR with narrative shocks to estimate the effects of standard monetary policy on the Gini coefficients for consumption and income. Instead, we focus on euro area QE, which is part of the non-standard monetary policy toolbox designed by central

[5] This calibration of the size of the QE shock to a 30 basis points drop in the term spread is close to the lower boundary estimated for the effect of the first QE announcement in the euro area, see for example Altavilla et al. (2015).

banks to circumvent the zero lower bound for nominal interest rates and we also assess the effects of monetary policy on wealth inequality. We find that the response of income inequality to QE in the euro area is qualitatively similar to that of income to standard policy in the US (as estimated by Coibion et al., 2017). In addition, we provide a decomposition of the effects on income into the extensive (i.e., the earnings heterogeneity channel) and the intensive (the income composition channel) margins.[6] By also looking at wealth inequality, we contribute to the debate on the relative importance of direct and indirect effects of monetary policy on consumption, since such effects can be estimated only by considering the transmission channels involving both income and wealth (Kaplan et al., 2018; Auclert, forthcoming). Casiraghi et al. (2018) (on Italian data) and, at least partly, Bunn et al. (2018) (on UK data) focus on unconventional monetary policy. For the analysis of the macroeconomic effects of the policies, Casiraghi et al. (2018) relies on the Banca d'Italia's assessment of two different "unconventional monetary policy" scenarios, with and without financial stress in the economy. Bunn et al. (2018), instead, is based on the evaluation of the effects of the conventional and unconventional monetary policy stimulus of Bank of England in 2008–2014 reported in Carney (2016). The two papers conclude that while households at the bottom of the income scale benefit more, the overall effect of monetary policy on income and wealth inequality has been rather small. Differently from these two papers, we focus specifically on Quantitative Easing, we take a multi-country approach and we estimate the effects of QE in a VAR for four euro area countries which, among other things, also takes into account the cross-country spillovers related to the monetary policy impulse.[7] Adam and Tzamourani (2016) quantify the effects of hypothetical scenarios on the evolution of various asset prices (stock, bond and house prices) focusing exclusively on the wealth of euro area households. Our analysis has a different focus from the work of Kuhn et al. (2017), who describe the *unconditional* historical evolution of the US wealth distribution, highlighting the contribution of house prices for the lower 90% of the households and of stock prices for the top 10%. Our purpose, instead, is to isolate the effects of quantitative easing on inequality and, for this reason, we use impulse responses from a VAR to identify the changes in the wealth distribution *conditional* on the effects of quantitative easing.

The remainder of the paper is organized as follows. Section 2 outlines our empirical method based on a multi-country VAR model and a simulation on household-level income and wealth data. Section 3 describes and interprets the empirical results and the main robustness checks. Section 4 concludes.

[6] A few papers follow in the steps of Coibion et al. (2017) for other countries. Mumtaz and Theophilopoulou (2017) provide similar evidence for the UK. Guerello (2018) finds that in the euro area standard expansionary monetary measures typically reduce the dispersion in the income distribution (in the data from the European Commission Consumer Survey). In aggregate panel data from 32 advanced and emerging market countries, Furceri et al. (2018) find that contractionary monetary policy shocks increase income inequality, on average. The effect is asymmetric—tightening of policy raises inequality more than easing lowers it—and depends on the state of the business cycle.

[7] In addition, our approach to distribute the aggregate impulse responses, which borrows from Ampudia et al. (2016), differs from how the households' responses, in particular the response of income components, are modelled in these papers: Bunn et al. (2018) do not model the transitions from unemployment to employment (the extensive margin) and Casiraghi et al. (2018) do not separate the earnings heterogeneity and the income composition channels.

2 Empirical Methodology

We estimate the effects of QE on wealth and income of individual households in two steps: First, we estimate a Bayesian VAR model on aggregate data and identify the effects of monetary policy shocks at the aggregate level. Second, we undertake a reduced-form simulation using micro data to distribute the aggregate effects on components of income and wealth across individual households. This section describes both steps in detail.

2.1 The BVAR Model and the Identification of Monetary Policy

We identify the effects of QE using a large vector autoregression (VAR) with country-specific variables for four large countries, euro area variables and US variables.[8] Such setup allows us to estimate possibly heterogeneous country responses to a common euro area QE shock. In more detail, to capture the dynamic interrelationships among the variables, we adopt the following VAR setting:

$$\begin{aligned} y_t &= C + B_1 y_{t-1} + \cdots + B_p y_{t-p} + \epsilon_t, \\ \epsilon_t &\sim \mathcal{N}(0, \Sigma), \end{aligned}$$

where y_t is an N-dimensional vector of time-series, $B_1, \ldots, B_p$ are $N \times N$ matrices of coefficients on the p lags of the variables, C is an N-dimensional vector of constants and Σ is the covariance matrix of the errors. The model is specified in terms of the annualized (log-)levels of the variables and, in our specification, we have $N = 25$ and $p = 5$. In particular, for each of the four countries (France, Germany, Italy and Spain) we consider real GDP, the GDP deflator, the unemployment rate, house prices and wages. We also include short- and long-term interest rates and stock prices for the euro area and real GDP and short-term rates for the US.[9] The variables are available at the quarterly frequency, for the sample 1999Q1 to 2016Q4.

Potentially, this model may be subject to the "curse of dimensionality" due to the large number of parameters to be estimated, relative to the available sample. In such circumstances, the estimation via classical techniques would very likely result in overfitting the data and large estimation uncertainty. De Mol et al. (2008) and Bańbura et al. (2010) showed that imposing informative priors which push the parameter values of the model toward those of naïve representations (as, for example, the random walk model) reduces estimation uncertainty without introducing substantial bias in the estimates, thanks to the tendency for most macroeconomic and financial variables to co-move. In fact, in presence of comovement, the information in the data strongly "conjures" against the prior and it allows the parameters to still reflect sample information even if very tight prior beliefs are enforced.

For this reason, we estimate the model with Bayesian techniques. The prior for the covariance matrix of the residuals Σ is Inverse-Wishart, while the prior for the

[8] See Appendix B for more details on the macroeconomic database, our estimation strategy and the identifying assumptions for the monetary policy shocks.

[9] The US GDP and short-term rates lead the euro area counterparts (see, for example, Giannone et al., 2010), which suggests that it is important to include them not to omit relevant information which could "pollute" our estimates of the structural shocks driving the fluctuations in the euro area.

autoregressive coefficients is (conditional on Σ) normal. As it is standard in the BVAR literature, we follow Litterman (1979) and parameterize the prior distribution to shrink the parameters toward those of the naïve and parsimonious random walk with drift model, $X_{i,t} = \delta_i + X_{i,t-1} + e_{i,t}$. Moreover, in order to address the tendency of VARs to overfit the data via their deterministic component (see Sims, 1996, 2000; Giannone et al., 2018, for an extensive discussion of this pathology of VARs), we also impose two priors on the sum of the VAR coefficients. The full specification and the estimation method used for the VAR model follows Giannone et al. (2015). The setting of the prior distributions depends on the hyperparameters which describe their informativeness for the model coefficients. For these parameters, we follow the theoretically grounded approach proposed by Giannone et al. (2015), which suggests to treat them as random, in the spirit of hierarchical modelling, and conduct posterior inference also on them. As hyper-priors (i.e., prior distributions for the hyperparameters), we use proper but almost flat distributions. For details on specification of the prior distribution see Appendix A.[10]

To estimate the effects of quantitative easing, we identify an exogenous asset purchase shock similarly to Baumeister and Benati (2013). In addition, we offset the response of the euro area policy interest rate via a series of standard monetary policy shocks. This scenario captures the fact that standard monetary did not react, over the course of the recent crises, to offset the effects of the asset purchases—instead, the policy rate remained at the (zero) lower bound. We identify the effects of asset purchases using a combination of zero and sign restrictions (employing the algorithm of Arias et al., 2018). The main identifying assumption is that an expansionary asset purchase shock decreases the term spread (defined as long-term minus short-term interest rate)[11] and has a positive impact on the real economy of the four countries under analysis. The decrease in the term spread on impact is entirely accounted for by the drop in the long-term interest rates, given that standard monetary policy (captured by the short-term interest rates) is assumed not to react on impact to the asset purchases. For what concerns the macroeconomic environment, we impose a positive sign on the responses of GDP. The responses of all other variables, i.e., the GDP deflator, the unemployment rate, wages and house prices in the four countries, the US variables and stock prices, are left unrestricted. Notice that all the identifying assumptions are only imposed on impact, i.e., for the same quarter in which the shock materializes. The standard monetary policy shock is identified via standard zero restrictions. In particular, we assume that a change in the short-term interest rate can only affect, on impact, the long-term interest rate and the stock prices.

[10]A few papers lend support to this strategy to model cross-country macroeconomic data, showing that VAR models of the type we adopt in this paper provide accurate out-of-sample forecasts of macroeconomic and financial variables in the euro area (see, for example, Angelini et al., 2018; Capolongo and Pacella, 2018). A similar framework has been also used to estimate the effects of common euro area monetary policy shocks on various countries by Altavilla et al. (2016) (for both standard monetary policy and outright monetary transactions, OMT) and Mandler et al. (2016) (for standard monetary policy shocks). To appropriately capture the transmission channels of QE to different components of household wealth and income, we add more variables such as house prices to the existing frameworks.

[11]The short-term rate is the 3-month Euribor; the long-term rate is the euro area 10-year government benchmark bond yield.

2.2 The Reduced-Form Simulation on Household-Level Wealth and Income Data

Table 1 provides a general overview of the methodology we adopt to distribute the aggregate effects estimated in the BVAR across individual households.

The analysis described in Table 1 is conducted using the second wave of the Household Finance and Consumption Survey (HFCS). The HFCS is a unique ex ante comparable household-level dataset, which contains rich information on the structure of income and household balance sheets and their variation across individual households. The dataset also collects information about socio-demographic variables, assets, liabilities, income and indicators of consumption. For most countries, the reference year of the HFCS wave 2 is 2014, which matches quite well the start of the Asset Purchase Programmes. We focus on the four largest euro area countries, in which the HFCS (net) sample ranges roughly between 4,500 households (Germany) and 12,000 households (France).[12] For Spain the reference year is 2011, for the other three countries 2014. To adequately capture the top tail of the distribution, wealthy households are over-sampled in most countries (including Spain, France and Germany).

2.2.1 Estimating the Effects of QE on Household Income: The Earnings Heterogeneity and the Income Composition Channels

Starting with our baseline characterization of the income composition channel which, in subsequent discussion we will also refer to as the intensive margin of QE, Figure 1 shows that the key income component for most households is income from employment and self-employment. We use impulse responses of wages to assess how these income components are affected by QE at the household level. For income from rental of properties, financial investments and pensions, instead, we assume that there is no change due to QE (Table 1). In section 3.2.3 we provide a robustness analysis to gauge the relevance of this no-change assumption for some categories of income such as, for example, financial income.

The earnings heterogeneity channel is instead related to the effect of monetary policy on employment. We model this extensive margin as follows. The aggregate results suggest that quantitative easing reduces the aggregate unemployment rate. In turn, household-level data on employment and income make it possible to simulate which unemployed people become employed and by how much their incomes increase. The simulation, which broadly follows the setup of Ampudia et al. (2016), is divided in two steps and runs at the individual level (not at the household level); the results are then aggregated to household level.

Step 1: Probit Simulation for the Employment Status

In the first step, we distribute the aggregate decline in unemployment across individuals, using a probit regression which takes into account individual characteristics. This allows

[12] See Household Finance and Consumption Network (2016), in particular Table 1.1, for information on the second wave of the HFCS.

us to pin-down which individuals become employed as a result of QE. More in details, for each country c, we first estimate a probit model regressing individual's i employment status Y on demographic characteristics:

$$\Pr(Y_i = 1|X_i = x_i) = \Phi(x_i'\beta_c), \tag{1}$$

where X denotes demographics: gender, education, age, marital status and the number of children; $\Phi(\cdot)$ denotes the normal cdf. For each individual we denote the fitted values, the estimated probability of being employed, as $\hat{Y}_{c,i}$ and we use it to simulate who becomes employed thanks to QE. This is done by drawing, for each person i, a uniformly distributed random 'employment' shock ξ_i. If the value of ξ_i is sufficiently below $\hat{Y}_{c,i}$ and the person is actually unemployed, she becomes employed. The threshold for moving into employment is computed to have a number of individuals becoming employed that is consistent with the VAR impulse response of the aggregate unemployment rate in each country.[13] We repeat the simulation many times and report the average results across repetitions.[14]

Step 2: Heckman Imputation of Labor Income

In the second step we replace unemployment benefits of people who are newly employed with wage, which is estimated based on their demographic characteristics. Technically, the wage of newly employed individuals is estimated by a two-step Heckman selection model. Our exclusion restrictions are the marital status and the presence of children. We assume these factors may affect the work status but not the wage of the employed. The remaining regressors in the model are gender, education and age.

2.2.2 Estimating the Effects of QE on Household Wealth: The Portfolio Composition Channel

To simulate the effects of quantitative easing on wealth, i.e., to capture the portfolio composition channel, we use the detailed *quantitative* information about holdings of various asset classes by each household in the HFCS (i.e., we know the nominal market value of each asset class owned by households). The effects of monetary policy on household wealth are obtained by multiplying the holding of each asset class (in EUR) by the corresponding change in asset prices given by the VAR impulse response.

In particular, our VAR includes three asset price variables: house prices, stock prices and bond prices. We multiply the holdings of housing wealth—i.e., household's main residence and other real estate—by house prices. We multiply the holdings of shares and household's self-employment businesses by stock prices.[15] Finally, we multiply the

[13] In practice, we sort unemployed individuals by their value of $(\xi_i - \hat{Y}_{c,i})$ and those with the lowest rank become employed until the reduction in the unemployment rate matches the value given by the impulse response. We use survey weights in this calculation.

[14] The empirical results in the paper are based on 200 iterations.

[15] As described in Table 1, we assume other classes of net wealth, most importantly deposits and liabilities remain unaffected by monetary policy. For the time period we focus on—since 2014—this seems reasonable as the short-run interest rate was at the zero lower bound. The HFCS also records holdings of voluntary pensions, for which we in the baseline scenario assume they are unaffected by stock prices. Data on Euro area insurance corporation and pension fund statistics, `https://www.ecb.europa.eu/press/pr/stats/icpf/html/index.en.html`, indicate that pension funds hold a

holdings of bonds by the change in the price of the 10-year bond implied by the initial decline in the long-term rate.

This calculation assumes that households do not adjust their portfolios in response to monetary policy. This assumption of no rebalancing seems a reasonable first-order approximation for two reasons. First, we consider responses to relatively *small* monetary policy shock over the short-run horizon of several quarters. Second, substantial evidence exists on the sluggishness in household portfolios. This holds not only for very illiquid assets (such as housing) but also for many financial assets. For example, a well-known paper by Ameriks and Zeldes (2004) documents that almost half of the households in their data on retirement accounts (held by TIAA–CREF) made no active changes to their portfolio of stock over the *nine-year* period they consider. Similar findings are reported in Bilias et al. (2010): The bulk of US households exhibit considerable inertia in their stock portfolios (held in brokerage accounts). Several papers examine inertia in household portfolios using high-quality administrative data. Fagereng et al. (2018) document evidence on the limited extent of rebalancing of illiquid and risky assets in response to receiving a lottery prize in Norwegian data. Using Danish data, Andersen et al. (2018) study the substantial inaction of households regarding mortgage refinancing. In Swedish data, Calvet et al. (2009) find very weak active rebalancing in the household sector as a whole, though at the household-level active rebalancing compensates about half of idiosyncratic passive variations in the risky share and is stronger for financially sophisticated households. In section 3.2.3 below, we also investigate how robust the results are to assuming some rebalancing in holdings of stocks and bonds.

3 Empirical Results

This section describes our estimates, first focusing on the effects of monetary policy on aggregate variables identified using the VAR model, then considering the effects on wealth and income of individual households via the three channels described in the previous section: (i) income composition, (ii) portfolio composition and (iii) earnings heterogeneity.[16]

3.1 Aggregate Effects of Quantitative Easing

We scale the size of the shock to a 30 basis point drop in the term spread. This normalization roughly matches the lower boundary of the estimated QE impacts on the term spread in existing studies on the euro area.[17] This normalization is imposed to offer a plausible quantification of the effects of QE on inequality.

small fraction of their assets in stocks, i.e., about 9% of total assets is held in equities (2016Q4). Notice however that 21.5% is held in investment funds, for which it is difficult to determine what fraction of their assets they hold in stocks.

[16] We do not consider other channels of transmission, such as the interest rate exposure channel of Auclert (forthcoming) and the inflation channel of Doepke and Schneider (2006). The former is analyzed quantitatively in Ampudia et al. (2018), while the latter turns out to have a negligible effect on inequality.

[17] For example, Altavilla et al. (2015) estimate that the first APP announcement in January 2015 was associated with a 30 to 50 basis point drop in the euro area ten-year sovereign bond yields.

Figure C.1 in Online Appendix C at the end of the paper reports all the impulse responses to the QE shock and the median response to the QE scenario in which the reaction of standard monetary policy to the QE shock is offset by standard monetary policy shocks. Our results are qualitatively in line with the previous literature, which finds relevant effects of asset purchases on the real economy—see Dell'Ariccia et al. (2018) (e.g., their Table 1) for an up-to-date overview of the literature. We also find that QE boosts the GDP deflator, wages and asset prices, although generally these results are surrounded by a larger uncertainty. To gauge the relevance of the effects of QE, notice that our quantitative easing shock (an exogenous drop by 30 basis points in the term spread) has roughly the same effect on GDP as a 100 basis point surprise drop in the policy rate.[18]

Figure 3 zooms on the median impulse responses of the variables that play an important role in our subsequent analysis on individual households. The term-spread shock has a relatively short-lived impact on the term spread[19] itself, whose median response is close to zero already after three quarters. The peak response of stock prices is quite large—4 %—but also quite transitory.

The country-specific impulse responses in Figure 3 document the extent of heterogeneity across the four countries. House prices increase in all countries but, for example, in Spain the increase is close to two percent, while in Germany it is about a third of that size. It is plausible that these differences in impulse responses arise due to different institutional settings. For example, Calza et al. (2013) show that house price responsiveness to monetary policy is significantly stronger in countries with larger flexibility/development of mortgage markets (e.g., in terms of the size of mortgage debt, extent of adjustable-rate mortgages or availability of equity release products; see also related work of Nocera and Roma, 2018). The responses of the labor market variables also display a marked heterogeneity across countries. The unemployment rates drop in all countries but, again, the response in Spain is about three times as large as in Germany, with Italy and France in between these two extremes. The response of wages, instead, also varies in sign, with a slight decrease in Spain and increases in other countries.

3.2 Effects of Quantitative Easing on Individual Households

We report the estimates of the effects on income and wealth of individual households using a series of figures with 'micro' impulse responses implied by the micro-simulation described in section 2.2. The impulse responses are grouped in terms of quintiles of the income and wealth distributions.

[18] Debortoli et al. (2018) estimate that standard monetary policy and quantitative easing work as perfect substitutes (in the US). Inoue and Rossi (2018) find that unconventional monetary policy has similar effects to conventional expansionary monetary policy, leading to an increase in both output growth and inflation.

[19] Notice that the long-term interest rate coincides with the term spread—given that the short-term interest rate is assumed not to change on impact to the QE shock, and that its response is zeroed out over the rest of the horizon by means of standard monetary policy shocks.

3.2.1 Effects on Household Income: The Earnings Heterogeneity and the Income Composition Channels

In the baseline setup, the effects of QE on income arise via two channels: (i) the earnings heterogeneity—the increase in income as people become employed (also defined as the extensive margin) and (ii) the income composition channel—the increase in labor income (for all employed people) due to higher wages (also defined as the intensive margin).

Let us first investigate the earnings heterogeneity channel in isolation. Figure 4 shows the impulse responses of the unemployment rate by (country-level) income quintiles. The first noteworthy result is that the stimulative effects on employment are strongly skewed toward low-income households. This finding is not straightforward because there are two countervailing factors that can affect the response of unemployment across income quintiles. On the one hand, higher income individuals have generally more favourable demographics (for example, an higher level of education) and, hence, also an higher estimated probability to become employed.[20] On the other hand, and this is the key factor to explain the result, the bottom right panel of Figure 4 shows that the number of unemployed is heavily skewed toward the bottom income quintile across all four countries.

Figure 4 also shows a relevant heterogeneity in the micro impulse responses across countries, both regarding the level and the dispersion of responses across income quintiles. One factor to explain the differences, in particular for the levels, is the cross-country difference in macro responses. For example, the overall reduction in unemployment is larger in Spain than in the other three countries. Instead, the dispersion of micro impulse responses across income quintiles is importantly affected by the distribution of the unemployed across quintiles, which is very different across countries. Indeed, a relevant mass of unemployed people in Spain has income in higher quintiles, so that the differences in impulse responses across quintiles in Spain are smaller (see, again, the bottom right panel in Figure 4). In contrast, the number of the unemployed in Germany and Italy is more strongly skewed toward the lowest income quintile, which causes unemployment in the lowest income quintile to drop more (relative to other quintiles) in these two countries. A final, although less relevant factor, that can explain the differences in the dispersion of micro responses is that the employment probabilities in the probit models (1) are country-specific.

Figure 5 shows the micro responses of mean income by income quintile, merging the earnings heterogeneity and the income composition channels. These responses are primarily driven by the transitions into employment and by differences in replacement rates (as estimated by the Heckman model). The replacement rates are in general more generous in Germany and France than in Spain and, in particular, Italy.[21] As a result, the magnitude and dispersion of income responses in Italy and Spain is larger. For

[20] In order to appreciate the quantitative relevance of this heterogeneity in probabilities to become employed, a counterfactual scenario where all individuals have the same probability to be drawn out of unemployment implies a significantly stronger stimulating effects on the lower income quintiles compared to our scenario based on estimated probabilities—as documented in Figure C.2 in the Online Appendix C. Obviously, although significant, this impact does not outweigh the countervailing effect due to over-representation of unemployed people in the lower income quintiles.

[21] See, e.g.,the OECD Statistics on Benefits and Wages: `http://www.oecd.org/els/benefits-and-wages-statistics.htm`.

example, the large positive response in mean income of the lowest quintile in Italy arises thanks to both the substantial decline in unemployment rate highlighted in Figure 4 and the substantial increase in (labor) income of the newly employed individuals.[22] These findings imply that the earnings heterogeneity channel is the most relevant to explain the changes in income across quintiles. To more precisely show this point, Figure 6 decomposes the overall increase in mean income into the extensive (earnings heterogeneity) and the intensive margins (income composition) for an aggregate of the four countries, one year after the shock. The extensive margin is particularly strong in the bottom income quintile, where wage growth plays a very small role. However, transitions from unemployment to employment make up the bulk of the total effect on income across much of the whole distribution (except for the top income quintile).

To summarize the effects of quantitative easing on income inequality, Table 2 shows that the Gini coefficient for gross household income declined from 43.07 to 42.86 when we isolate the effects of QE, one year after the shock.

3.2.2 Effects on Household Wealth: The Portfolio Composition Channel

This section analyses how the portfolio composition channel affects household net wealth. Figure 7 shows the micro responses of median net wealth by wealth quintile.[23] These responses arise from a combination of the response of house prices, stock prices and bond prices, and holdings of wealth across the distribution (and countries). Broadly, the responses of wealth in quintiles two to five increase by around 1.5% in France, Spain and Italy, and are rather flat in Germany. There is little evidence that the median wealth among the top wealth quintile households would increase more strongly, though this does happen for the top 10% of the wealth distribution, where the holdings of stocks are prevalent. Overall, Table 2 documents that the Gini coefficient on net wealth is only modestly affected by QE, one year after the shock. An important takeaway from this exercise is the key role of including house prices in the analysis, since most households own large holdings of housing wealth rather than stocks and bonds, which are only relatively more prominent in the top tail of the distribution.[24]

3.2.3 Robustness Checks

This section explores whether some plausible perturbations of our baseline specification affect the main results. For these robustness checks, we rely on alternative macroeconomic data sources as, for example, the data on the flow of funds of the four countries under analysis. In order to derive the effects of the QE scenario on these variables, we extract the time-series of the QE shocks from our multi-country VAR, and we compute

[22]The results are shown for gross (pre-tax) income. The increase in after-tax income would be somewhat lower, however, not by much, as most newly employed people are not subject to large taxes. As for the effect on inequality of net income, it would be reduced more than inequality of gross income because of progressivity of taxes.

[23]The growth rate for the lowest quintile is not shown because its level is close to EUR 0.

[24]This finding is in line with Adam and Tzamourani (2016); see, e.g., their Figure 4. See also Kuhn et al. (2017), Figure 17 for historical evidence from the US.

the responses of the variables to the shocks by means of the local linear projection method of Jordà (2005).[25]

Our baseline analysis on the effects of QE on the income distribution via the income composition channel neglects the effects of QE on financial income[26] and, hence, it might bias our analysis of the impact of QE on income inequality, to the extent that financial income particularly characterizes the top tail of the income distribution. Indeed, if quantitative easing increases financial income, e.g., via stimulating corporate profits, this effect may to some extent work counter to the employment effect and widen income inequality.[27] To address this concern, we use the local linear projection method to investigate how two alternative (aggregate) measures of financial income respond to quantitative easing: (i) profits (available for the euro area) and (ii) net property income (available for the four country under analysis). Figure C.3 in Online Appendix C shows that profits increase by up to about 5% (despite substantial estimation uncertainty), one year after the shock. Figure C.4, instead, shows a relevant heterogeneity across countries in the reaction of net property income to the QE shock, one year after the shock, with the smallest increase in Germany (about 4%) and the largest in Italy (about 20%). Figure 8 considers the implications for the income distribution: (i) assuming that financial income behaves similarly to profits (i.e., also increasing by 5 % in all countries), top panel; (ii) assuming that financial income responds as estimated by linear projections for aggregate data on net property income, bottom panel. As expected, the scenarios increase income in particular among the top income quintile of households. However, the overall impact on total income is quite limited and, as shown in Table 2, the Gini coefficients for the two scenarios only marginally change with respect to our baseline assessment.

Turning to wealth inequality, first we relax our assumption of no portfolio rebalancing. To get an idea of a plausible amount of rebalancing, we rely on country-level flow-of-funds data on the holdings of different asset categories by households. As a caveat to this analysis, notice that the data refer to the value of holdings and, hence, they also reflect asset valuations which, as we have seen, are affected by QE. This could be a source of mis-measurement for the impact of QE on the volume of asset holdings and, hence, these results should be only taken as suggestive. As for the analysis on income, we estimate how quantitative easing affects holdings of wealth components using local linear projections, and we find that QE affects mostly the value of stock holdings. In fact, despite substantial estimation uncertainty, Figure C.5 in Online Appendix C suggests quite large increases (in the median responses) in the holdings of shares. Figure 9 takes these estimates to micro data, showing a scenario in which households *buy* 15% of their holdings of stocks in response to quantitative easing. We find that stock trading affects the distribution of net wealth only very little: in particular, Table 2 documents that the

[25]See Appendixes A and B for the description of the local linear projection method and for more information on the alternative data sources we use in the robustness checks.

[26]Financial income includes income in the form of interest or dividends on sight deposits, time and saving deposits, certificates of deposit, managed accounts, bonds, publicly traded stock shares or mutual funds. More broadly, we also include income from renting real estate and income from private business other than self-employment.

[27]Existing evidence, e.g., Guvenen et al. (2014), points to slight, rather than strong, pro-cyclicality in the unconditional dynamics of earnings and financial income among top earners.

Gini coefficient on wealth under this alternative scenario falls to 68.08, one year after the shock, rather than to the value of 68.04 which we had found in our baseline scenario. This is explained by the fact that the share of stocks in the portfolios of European households lies below 5%. We view this finding as an upper bound on the how active portfolio rebalancing can affect wealth inequality because evidence from micro data, including the influential work of Calvet et al. (2009) and Brunnermeier and Nagel (2008), typically estimates that (if at all) individual households tend to actively rebalance in the opposite direction, i.e., by selling risky financial assets after experiencing high returns.

Another important aspect that our baseline scenario could be disregarding, considering the relevance of housing for the wealth distribution in Europe, is the potential heterogeneity in the responses of house prices across regions (arising, e.g., due to differences in elasticity of housing supply). To investigate the relevance of such scenario, Figure C.6 shows the dispersion in responses of house prices across provinces in Spain,[28] confirming some, though not overwhelming heterogeneity: a 68% confidence range around the aggregate response after 4 quarters spans increases between 0 and 6% in local house prices. Interestingly, Figure C.7 documents that the percentage increase in house prices tends to be larger in provinces with higher levels of house prices (measured in EUR per square meter), so that more expensive houses respond more strongly to monetary policy. To assess how this heterogeneity in responses of house prices to monetary policy affects our baseline results on the portfolio composition channel, we undertake the following simulation. The HFCS dataset collects information both on the price and on the area of the household main residence (in square meters). Within each of the four countries, we sort households into five quintiles by the price per square meter. In line with the scatter plot in Figure C.7, we then assume that quantitative easing increases the prices of more expensive houses (in terms of the price per square meter) more strongly.[29] Figure 10 quantifies how our baseline compares to the simulation in which the increase in house prices depends on the level of house prices. Because poorer households tend to own less expensive houses, the alternative assumption reduces the differences in changes in wealth across quintiles: For the lowest net wealth quintile, median wealth grows by 1.8% (compared to 2.4% for the baseline)—still quite a bit above the change for the other quintiles (which remains around 1%). Table 2 shows that under this scenario the Gini coefficient on net wealth remains essentially unchanged at 68.09 and, hence, our conclusion about the negligible effect of quantitative easing on wealth inequality remains unaffected.

[28] Spain is the only country in our sample for which quarterly data on regional house prices are available since 1999.

[29] Specifically, we calibrate that across the quintiles, the responses of the price of the household main residence and its other real estate after 4 quarters range between 0–4% for Spain, between 0–3% for France and Italy and between 0–1% for Germany. This calibration thus preserves the aggregate response of house prices to quantitative easing estimated in the VAR, upper right-hand panel in Figure 3, and adds to it a positive relationship between the level of house prices and their sensitivity to monetary policy.

4 Conclusions

Combining estimates from a VAR with aggregate data and a simulation on household-level data, we quantify how the recent quantitative easing measures in the euro area affect individual households via the portfolio composition, the income composition and the earnings heterogeneity channels. We find that although QE has only negligible effects on wealth inequality, it noticeably compresses the income distribution since many households with lower incomes become employed. Specifically, a year after the shock, the Gini coefficient for income falls from 43.1 to 42.9, while the reduction of the Gini coefficient for net wealth is an order of magnitude smaller.

The effects of monetary policy are likely to fade away in the long run and, hence, QE should not be a key driver of inequality in the long run, when other factors, such as globalization or progressivity of the tax system are more important. However, our results suggest that quantitative easing substantially contributed to support vulnerable households.

Our results are also informative about the strength and nature of the transmission of monetary policy to consumption. An extensive literature has recently documented that constrained households—e.g., those with low incomes or little liquid assets—have high marginal propensities to consume. We find such households also particularly benefit from a monetary stimulus, which boosts their employment and income. In combination, these two facts imply that the stimulating effect of quantitative easing on aggregate consumption is substantially magnified both because it disproportionately boosts incomes in the lower part of the distribution and because this impulse has a stronger effect on consumption via the larger MPCs of the constrained households.[30]

[30] Ampudia et al. (2018) quantify the channels of monetary transmission to consumption and their heterogeneity across households.

Appendix A: Estimation

A.1 The Prior Distributions

The prior distributions in our Bayesian VAR are specified as follows. For the prior on the covariance matrix of the errors, we set the degrees of freedom of the Inverse-Wishart distribution equal to $N+2$, the minimum value that guarantees the existence of the prior mean, and we assume a diagonal scaling matrix Ψ, which we treat as a hyperparameter.

The baseline prior on the model coefficients is a version of the Minnesota prior (see Litterman (1979)). This prior is centered on the assumption that each variable follows an independent random walk process, possibly with drift. The prior first and second moments for the VAR coefficients are:

$$\begin{aligned} \mathbf{E}\big((B_s)_{ij}\big|\Sigma\big) &= \begin{cases} 1 & \text{if } i=j \text{ and } s=1 \\ 0 & \text{otherwise} \end{cases}, \\ \operatorname{cov}\big(\,(B_s)_{ij}\,,(B_r)_{hm}\,\big|\Sigma\big) &= \begin{cases} \lambda^2 \frac{1}{s^2} \frac{\Sigma_{ih}}{\psi_j/(d-n-1)} & \text{if } m=j \text{ and } r=s \\ 0 & \text{otherwise} \end{cases}. \end{aligned}$$

Notice that the variance of this prior is lower for the coefficients associated with more distant lags and that coefficients associated with the same variable and lag in different equations are allowed to be correlated. Finally, the key hyperparameter is λ, which controls the scale of all variances and covariances and effectively determines the overall tightness of this prior. The terms Σ_{ih}/Ψ_j account for the relative scale of the variables. The prior for the intercept C is non-informative.

The Minnesota prior is complemented with two priors on the sum of the VAR coefficients, introduced as refinements of the Minnesota prior to further "favor unit roots and cointegration, which fits the beliefs reflected in the practices of many applied macroeconomists" (see Sims and Zha (1998), p. 958). These additional priors tend to reduce the importance of the deterministic component implied by VARs estimated conditioning on the initial observations (see Sims (1996) and Giannone et al. (2015)). The first of these two priors is known as no-cointegration (or, simply, sum-of-coefficients) prior.

To understand what this prior entails, we rewrite the VAR equation in an error-correction form:

$$\Delta y_t = C + (B_1 + \cdots + B_p - I_N)y_{t-p} + A_1\Delta y_{t-1} + \cdots + A_p\Delta y_{t-p} + \epsilon_t,$$

where $A_s = -B_{s+1} - \cdots - B_p$. A VAR in first differences implies the restriction $\Pi = (B_1 + \cdots + B_p - I_N) = 0$. Doan et al. (1984) introduced the no-cointegration prior which centered at 1 the sum of coefficients on own lags for each variable, and at 0 the sum of coefficients on other variables' lags. This prior also introduces correlation among the coefficients on each variable in each equation. The tightness of this additional prior is controlled by the hyperparameter μ. As μ goes to infinity, the prior becomes diffuse, while as it goes to 0, it implies the presence of a unit root in each equation.

The fact that, in the limit, the prior just discussed is not consistent with cointegration motivates the use of an additional prior on the sum of coefficients that was introduced by

Sims (1996) and is known as dummy-initial-observation prior. This prior states that a no-change forecast for all variables is a good forecast at the beginning of the sample. The hyperparameter δ controls the tightness of this prior. As δ tends to 0, the prior becomes more dogmatic and all the variables of the VAR are forced to be at their unconditional mean, or the system is characterized by the presence of an unspecified number of unit roots without drift. As such, the dummy-initial observation prior is consistent with cointegration.

The setting of the prior distributions depends on the hyperparameters λ, μ, δ and Ψ, which describe the informativeness of the prior distributions for the model coefficients. In setting these parameters, we follow the theoretically grounded approach proposed by Giannone et al. (2015), who suggest to treat the hyperparameters as additional parameters, in the spirit of hierarchical modelling. As hyper-priors (i.e., prior distributions for the hyperparameters), we use proper but almost flat distributions.

A.2 Local Linear Projection

Our robustness exercises in section 3.2.3 adopt the local linear projection to derive the response of various variables to the shocks we estimate in the VAR. Let us briefly describe our application of the method developed in Jordà (2005). Denote X_t an additional variable of interest. We transform these variables as for the VAR, i.e., we compute annualized log-levels unless the variable is already expressed in terms of rates. Denote x_t the transformed variable.

Denoting u_t the time series of the QE structural shock derived from our VAR, we evaluate the impulse response β^h of x_t to the shock u_t at the horizon h by regressing x_{t+h} on u_t and the lags of x_t, i.e., we estimate the following regression:

$$x_{t+h} = \alpha + \beta^h u_t + \gamma(L)x_t + \varepsilon_t.$$

The regression is estimated by means of Bayesian techniques. We impose a flat prior on α and β^h, while we impose an informative prior on the coefficients on the lags, $\gamma(L)$. The informative prior has the exact same features of the Minnesota prior described in Appendix A. Notably, the shrinkage of the lagged terms grows with the horizon h at which the impulse response is computed.

Also for the local linear projections, we aim to evaluate the effects of the “QE scenario" in which standard monetary policy does not react to stabilize the economy. Hence, as in the VAR analysis, we estimate the response of all the alternative variables by means of linear projections on the VAR standard monetary policy shock. Then, we use these local linear projections to eliminate the effects of the response of the euro area policy interest rate from the local linear projection to the QE shock, using the same series of standard monetary policy shocks used in the VAR analysis.

Appendix B: Macroeconomic Data and Identification Assumptions

Table 3 describes our aggregate time series and sign restrictions we use to identify the effects of quantitative easing in our VAR.

In our robustness exercises, we exploit some additional data sources, available at the quarterly frequency in the sample 1999Q1–2016Q4. First, we consider data on profits for the euro area; precisely, this variable captures gross operating surplus (total economy, nominal, seasonally adjusted data) and is available from the Main National Accounts collection in the ECB Statistical Data Warehouse (SDW). The data on net property income and stock holdings of the four countries under analysis come from the Euro Area Sectoral Accounts. Finally, the data on regional house prices in Spain are available from the website of the Spanish government, Ministerio de Fomento.[31]

[31] We use the series "valor tasado medio de vivienda libre" (the aggregate house price, total national, and the house prices of the 17 regions for which the quarterly data are available, i.e., we exclude the autonomous cities Ceuta and Melilla): http://www.fomento.gob.es/BE2/?nivel=2&orden=35000000.

References

Acemoglu, Daron (2002), "Technical Change, Inequality, and the Labor Market," *Journal of Economic Literature*, 40(1), 7–72.

Adam, Klaus, and Panagiota Tzamourani (2016), "Distributional Consequences of Asset Price Inflation in the Euro Area," *European Economic Review*, 89(C), 172–192.

Altavilla, Carlo, Giacomo Carboni, and Roberto Motto (2015), "Asset Purchase Programmes and Financial Markets: Lessons from the Euro Area," working paper 1864, European Central Bank.

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.

Alvaredo, Facundo, Anthony B. Atkinson, Thomas Piketty, and Emmanuel Saez (2013), "The Top 1 Percent in International and Historical Perspective," *Journal of Economic Perspectives*, 27(3), 3–20.

Ameriks, John, and Stephen P. Zeldes (2004), "How Do Household Portfolio Shares Vary with Age?" working paper 9, Columbia University.

Ampudia, Miguel, Dimitris Georgarakos, Jiri Slacalek, Oreste Tristani, Philip Vermeulen, and Giovanni L. Violante (2018), "Monetary Policy and Household Inequality," working paper 2170, European Central Bank.

Ampudia, Miguel, Akmaral Pavlickova, Jiri Slacalek, and Edgar Vogel (2016), "Household Heterogeneity in the Euro Area since the Onset of the Great Recession," *Journal of Policy Modeling*, 38(1), 181–197.

Andersen, Steffen, John Y. Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai (2018), "Sources of Inaction in Household Finance: Evidence from the Danish Mortgage Market," mimeo.

Andrade, Philippe, Johannes Breckenfelder, Fiorella De Fiore, Peter Karadi, and Oreste Tristani (2016), "The ECB's Asset Purchase Programme: An Early Assessment," working paper 1956, European Central Bank.

Angelini, Elena, Magdalena Lalik, Michele Lenza, and Joan Paredes (2018), "Mind the Gap: A Multi-Country BVAR Benchmark for the Eurosystem Projections," mimeo, European Central Bank.

Arias, Jonas E., Juan Francisco Rubio-Ramírez, and Daniel F. Waggoner (2018), "Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications," *Econometrica*, 86(2), 685–720.

Auclert, Adrien (forthcoming), "Monetary Policy and the Redistribution Channel," *American Economic Review*.

Autor, David H. (2014), "Skills, Education and the Rise of Earnings Inequality among the 'Other 99 Percent'," *Science*, 344(6186), 843–851.

Bańbura, Marta, Domenico Giannone, and Lucrezia Reichlin (2010), "Large Bayesian Vector Autoregressions," *Journal of Applied Econometrics*, 25(1), 71–92.

Baumeister, Christiane, and Luca Benati (2013), "Unconventional Monetary Policy and the Great Recession: Estimating the Macroeconomic Effects of a Spread Compression at the Zero Lower Bound," *International Journal of Central Banking*, 9(2), 165–212.

Bilias, Yannis, Dimitris Georgarakos, and Michael Haliassos (2010), "Portfolio Inertia and Stock Market Fluctuations," *Journal of Money, Credit and Banking*, 42(4), 715–742.

Brunnermeier, Markus K., and Stefan Nagel (2008), "Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-evidence on Individuals," *American Economic Review*, 98(3), 713–736.

Bunn, Philip, Alice Pugh, and Chris Yeates (2018), "The Distributional Impact of Monetary Policy Easing in the UK Between 2008 and 2014," staff working paper 720, Bank of England.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini (2009), "Fight or Flight? Portfolio Rebalancing by Individual Investors," *The Quarterly Journal of Economics*, 124(1), 301–348.

Calza, Alessandro, Tommaso Monacelli, and Livio Stracca (2013), "Housing Finance and Monetary Policy," *Journal of the European Economic Association*, 11, 101–122.

Capolongo, Angela, and Claudia Pacella (2018), "Forecasting Inflation in the Euro Area: Countries Matter!" mimeo, Université Libre de Bruxelles.

Carney, Mark (2016), "The Spectre of Monetarism," speech, Liverpool John Moores University. https://www.bankofengland.co.uk/speech/2016/the-spectre-of-monetarism

Casiraghi, Marco, Eugenio Gaiotti, Lisa Rodano, and Alessandro Secchi (2018), "A "Reverse Robin Hood"? The Distributional Implications of Non-standard Monetary Policy for Italian Households," *Journal of International Money and Finance*, forthcoming.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia (2017), "Innocent Bystanders? Monetary Policy and Inequality," *Journal of Monetary Economics*, 88(C), 70–89.

Colciago, Andrea, Anna Samarina, and Jakob de Haan (2018), "Central Bank Policies and Income and Wealth Inequality: A Survey," working paper 594, Netherlands Central Bank, Research Department.

De Mol, Christine, Domenico Giannone, and Lucrezia Reichlin (2008), "Forecasting Using a Large Number of Predictors: Is Bayesian Shrinkage a Valid Alternative to Principal Components?" *Journal of Econometrics*, 146(2), 318–328.

Debortoli, Davide, Jordi Galí and Luca Gambetti (2018), "On the Empirical (Ir)Relevance of the Zero Lower Bound Constraint," working paper 1013, Barcelona Graduate School of Economics.

Dell'Ariccia, Giovanni, Pau Rabanal, and Damiano Sandri (2018), "Unconventional Monetary Policies in the Euro Area, Japan, and the United Kingdom," *Journal of Economic Perspectives*, 32(4), 147–172.

Doan, Thomas, Robert B. Litterman, and Christopher A. Sims (1984), "Forecasting and Conditional Projection Using Realistic Prior Distributions," *Econometric Reviews*, 3(1), 1–100.

Doepke, Matthias, and Martin Schneider (2006), "Inflation and the Redistribution of Nominal Wealth," *Journal of Political Economy*, 114(6), 1069–1097.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri (2016), "Heterogeneity and Persistence in Returns to Wealth," working paper 22822, National Bureau of Economic Research.

Fagereng, Andreas, Martin B. Holm, and Gisle James Natvik (2018), "MPC Heterogeneity and Household Balance Sheets," CESifo Working Paper Series 7134, CESifo Group Munich.

Furceri, Davide, Prakash Loungani, and Aleksandra Zdzienicka (2018), "The Effects of Monetary Policy Shocks on Inequality," *Journal of International Money and Finance*, 85(C), 168–186.

Giannone, Domenico, Michele Lenza, and Giorgio Primiceri (2015), "Prior Selection for Vector Autoregressions," *The Review of Economics and Statistics*, 97(2), 436–451.

Giannone, Domenico, Michele Lenza, and Giorgio Primiceri (2018), "Priors for the Long Run," *Journal of the American Statistical Association*, (forthcoming).

Giannone, Domenico, Michele Lenza, and Lucrezia Reichlin (2010), "Business Cycles in the Euro Area," in *Europe and the Euro*, 141–167, University of Chicago Press. https://ideas.repec.org/h/nbr/nberch/11669.html

Guerello, Chiara (2018), "Conventional and Unconventional Monetary Policy vs. Households Income Distribution: An Empirical Analysis for the Euro Area," *Journal of International Money and Finance*, forthcoming.

Guvenen, Fatih, Greg Kaplan, and Jae Song (2014), "How Risky Are Recessions for Top Earners?" *American Economic Review*, 104(5), 148–53.

Household Finance and Consumption Network (2016), "The Household Finance and Consumption Survey—Results from the Second Wave," Statistics Paper Series 18, European Central Bank, https://www.ecb.europa.eu/pub/pdf/scpsps/ecbsp18.en.pdf.

Hubmer, Joachim, Per Krusell, and Anthony A. Smith, Jr. (2018), "A Comprehensive Quantitative Theory of the U.S. Wealth Distribution," mimeo, Yale University.

Inoue, Atsushi, and Barbara Rossi (2018), "The Effects of Conventional and Unconventional Monetary Policy: A New Approach," mimeo, Universitat Pompeu Fabra.

Jordà, Òscar (2005), "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95(1), 161–182.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante (2018), "Monetary Policy Accoring to HANK," *American Economic Review*, 108(3), 697–743.

Katz, Lawrence F., and David H. Autor (1999), "Changes in the Wage Structure and Earnings Inequality," in Orley Ashenfelter and David Card, editors, *Handbook of Labor Economics*, volume 3 of *Handbook of Labor Economics*, chapter 26, 1463–1555, Elsevier.

Katz, Lawrence F., and Kevin M. Murphy (1992), "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 107(1), 35–78.

Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins (2017), "Income and Wealth Inequality in America, 1949–2016," discussion paper 12218, Centre for Economic Policy Research.

Kuznets, Simon (1955), "Economic Growth and Income Inequality," *American Economic Review*, 45, 1–28.

Litterman, Robert (1979), "Techniques of Forecasting Using Vector Autoregressions," working paper 115, Federal Reserve of Minneapolis.

Mandler, Martin, Michael Scharnagl, and Ute Volz (2016), "Heterogeneity in Euro-Area Monetary Policy Transmission: Results from a Large Multi-Country BVAR Model," discussion paper 03, Deutsche Bundesbank.

Mumtaz, Haroon, and Angeliki Theophilopoulou (2017), "The Impact of Monetary Policy on Inequality in the UK. An Empirical Analysis," *European Economic Review*, 98(C), 410–423.

Nocera, Andrea, and Moreno Roma (2018), "House Prices and Monetary Policy in the Euro Area: Evidence from Structural VARs," working paper 18-13, USC Dornsife INET.

Piketty, Thomas (2013), *Capital in the Twenty-First Century*, Harvard University Press.

Sims, Christopher A. (1996), "Inference for Multivariate Time Series Models with Trend," mimeo, Princeton University.

Sims, Christopher A. (2000), "Using a Likelihood Perspective to Sharpen Econometric Discourse: Three Examples," *Journal of Econometrics*, 39, 949–968.

Sims, Christopher A., and Tao Zha (1998), "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 39(4), 949–968.

Figure 1 Composition of Income

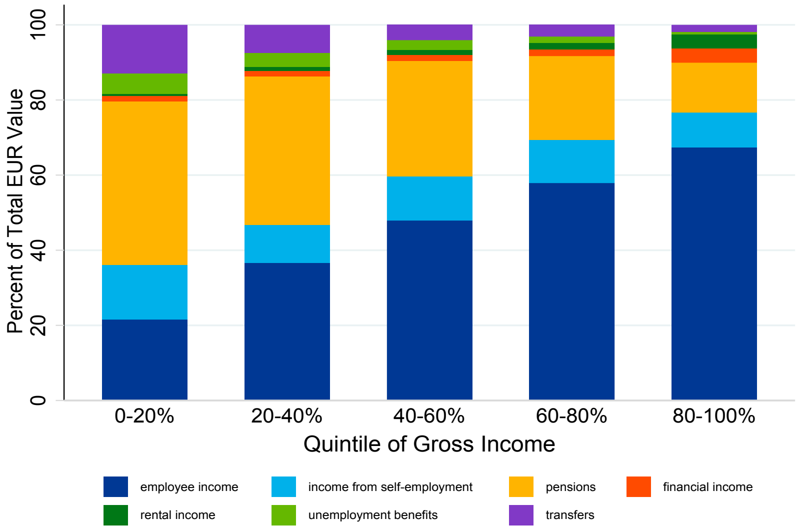

Source: Household Finance and Consumption Survey
Note: The chart shows how the share of income components in total gross income varies across quintiles of gross income. Unemployment benefits and transfers include regular social transfers (except pensions) and private transfers. The chart covers the euro area countries and includes the 17 countries included in wave 2 of the HFCS.

Figure 2 Composition of Total Assets

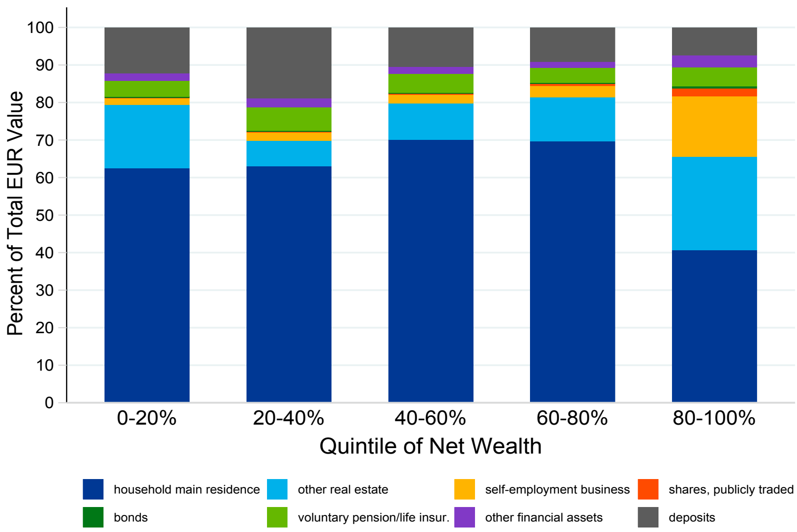

Source: Household Finance and Consumption Survey
Note: The chart shows how the share of components in total assets varies across quintiles of net wealth. Other financial assets include managed accounts, mutual funds and money owed to households. The chart covers the euro area countries and includes the 17 countries included in wave 2 of the HFCS.

Figure 3 Quantitative Easing Scenario: VAR Impulse Responses

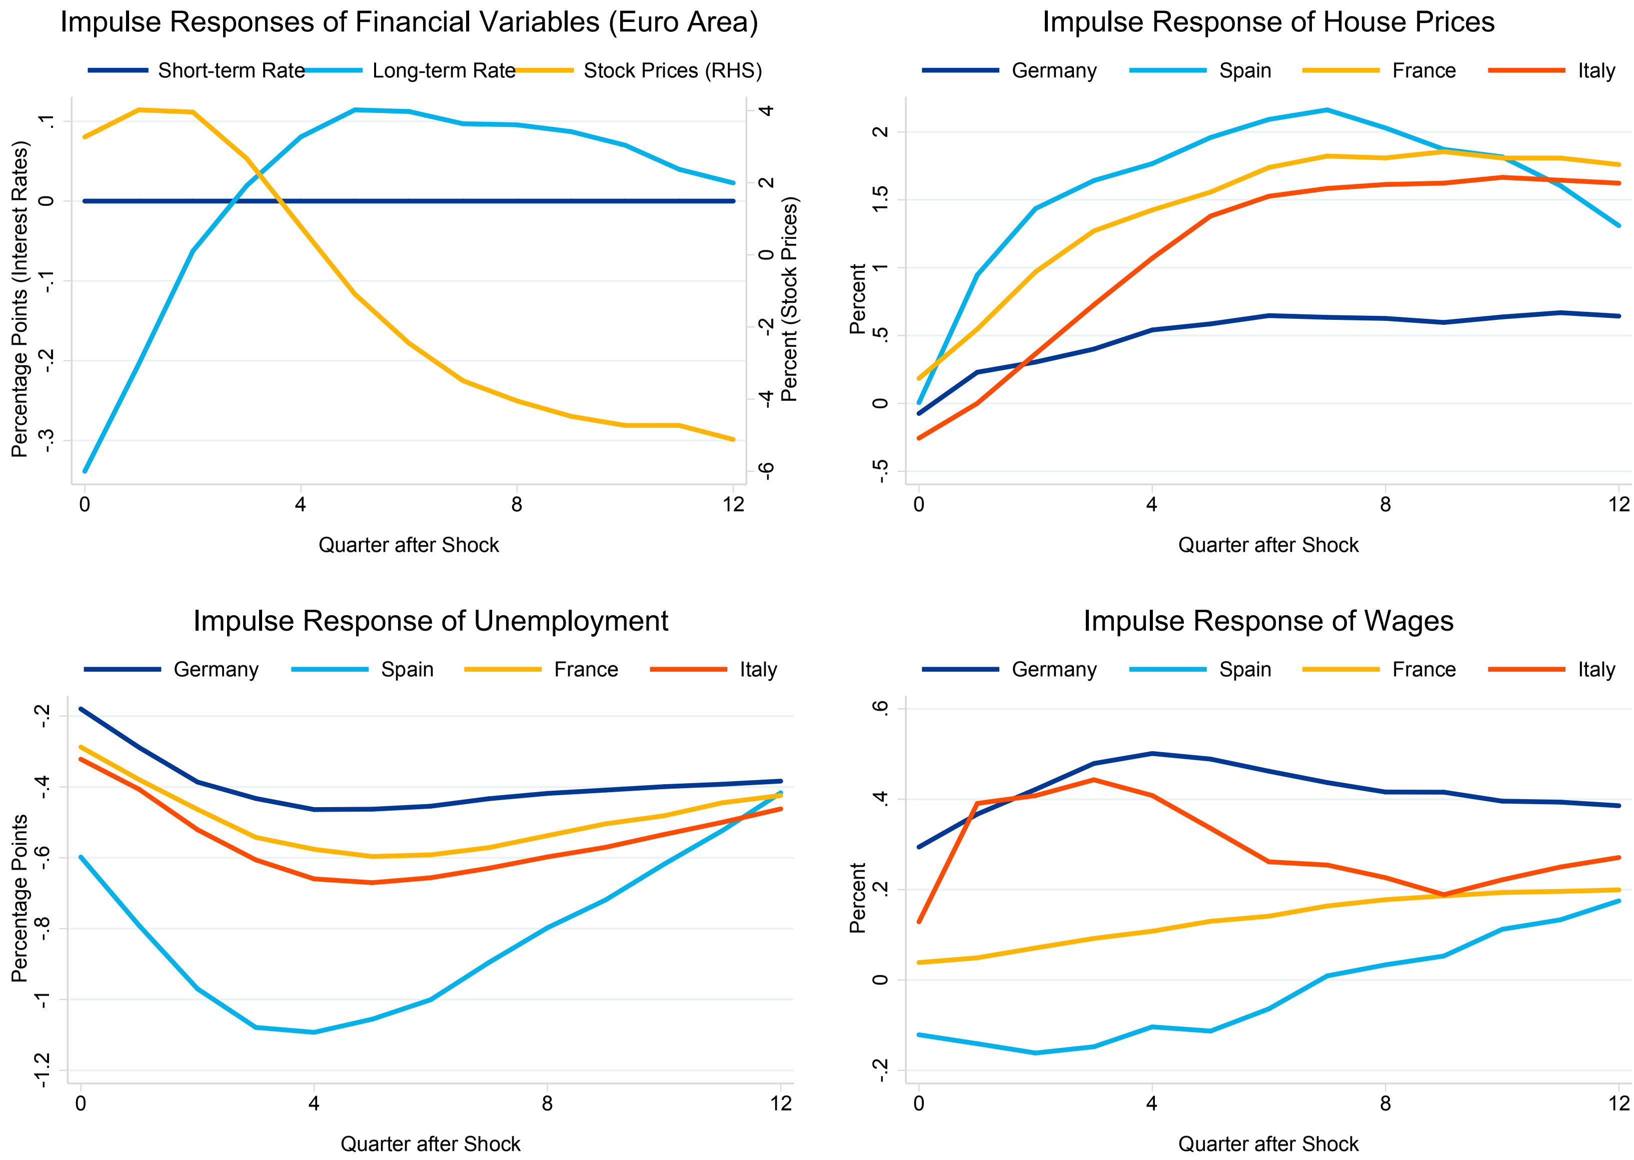

Note: The charts plots median responses to the asset purchase scenario. House prices: percentage deviation from baseline levels; wages: percentage deviation from baseline levels; unemployment rate: deviation from baseline level; stock price: percentage deviation from baseline levels.

Figure 4 Impulse Responses of Unemployment Rate by Country and Income Quintile

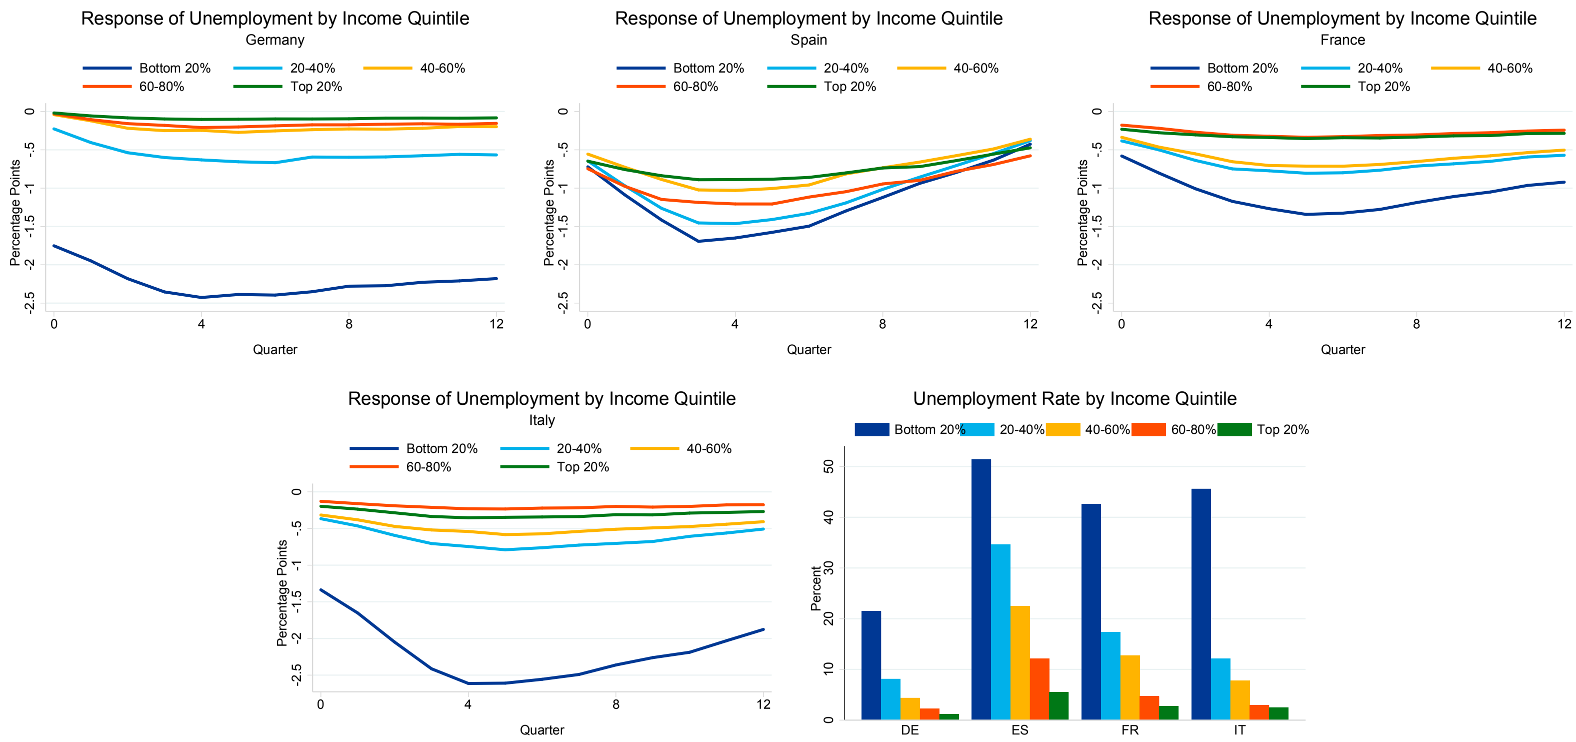

Source: Household Finance and Consumption Survey
Note: The charts show impulse responses of unemployment by income quintile.

Figure 5 Impulse Responses of Mean Income by Country and Income Quintile

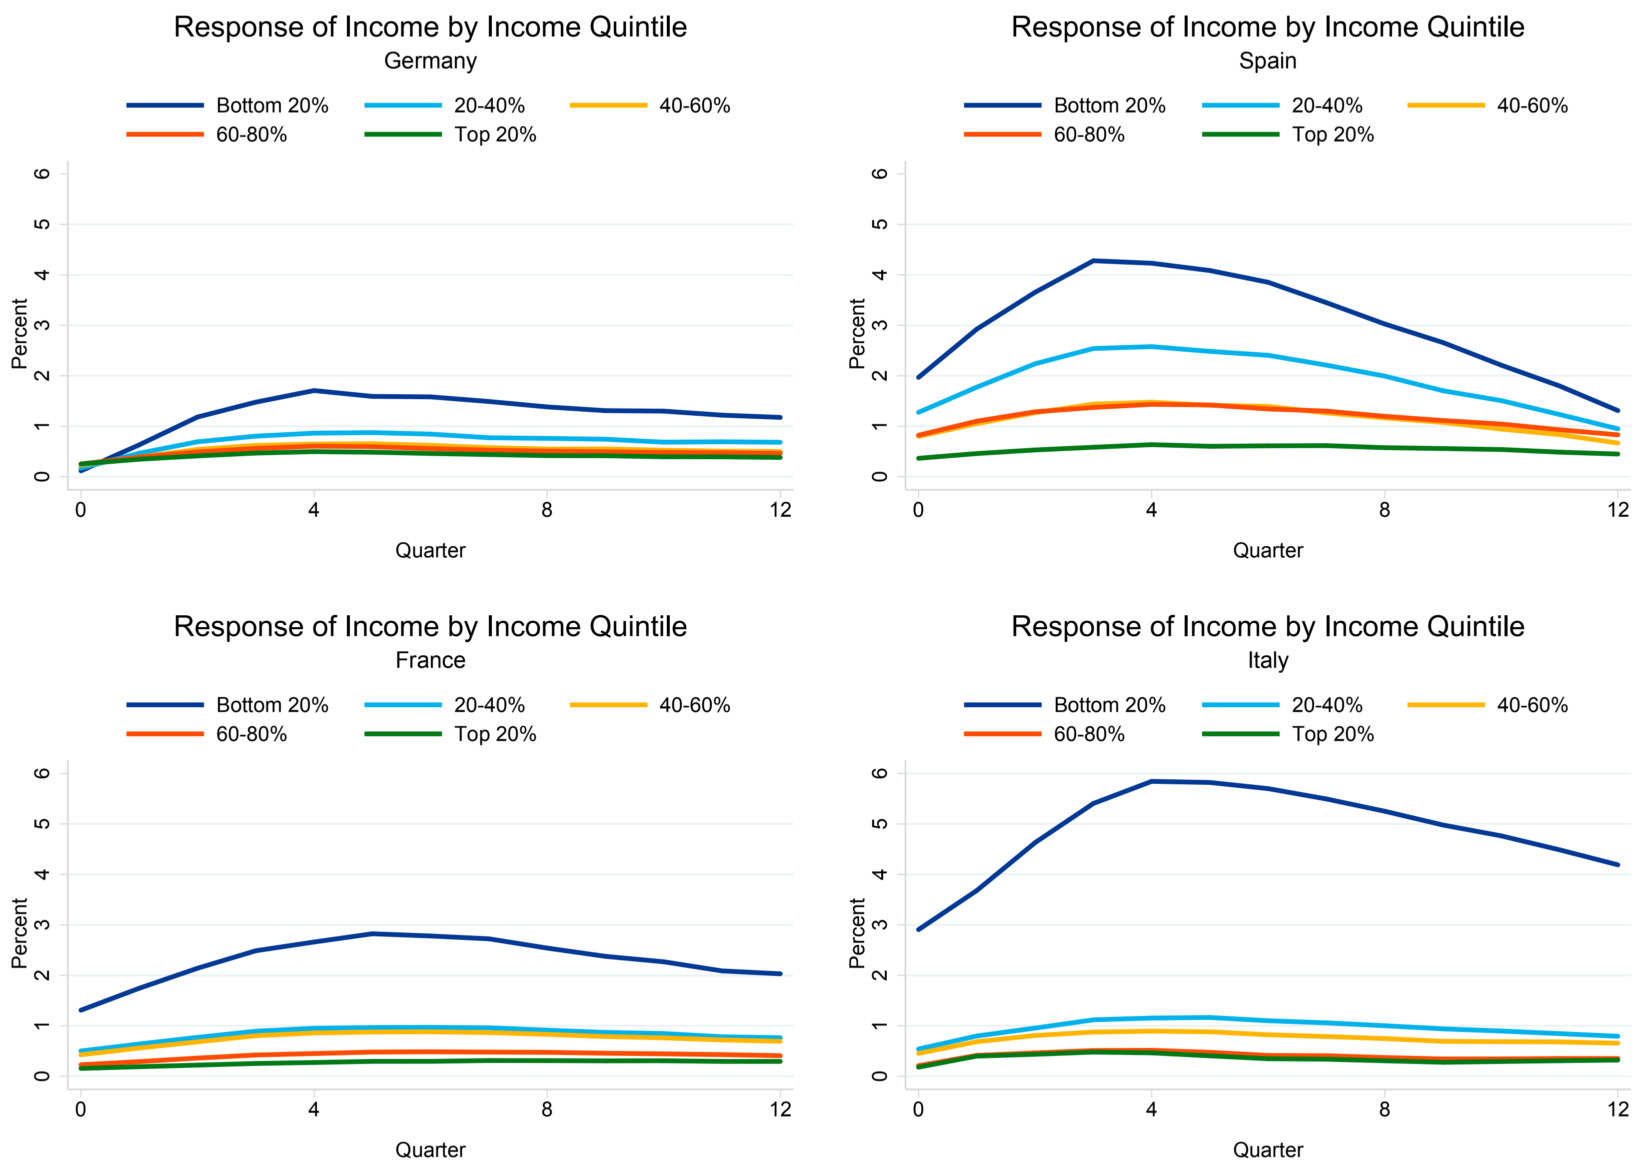

Source: Household Finance and Consumption Survey
Note: The charts show impulse responses of mean income by income quintile.

Figure 6 Decomposition of the Total Effect on Mean Income into the Extensive and the Intensive Margin

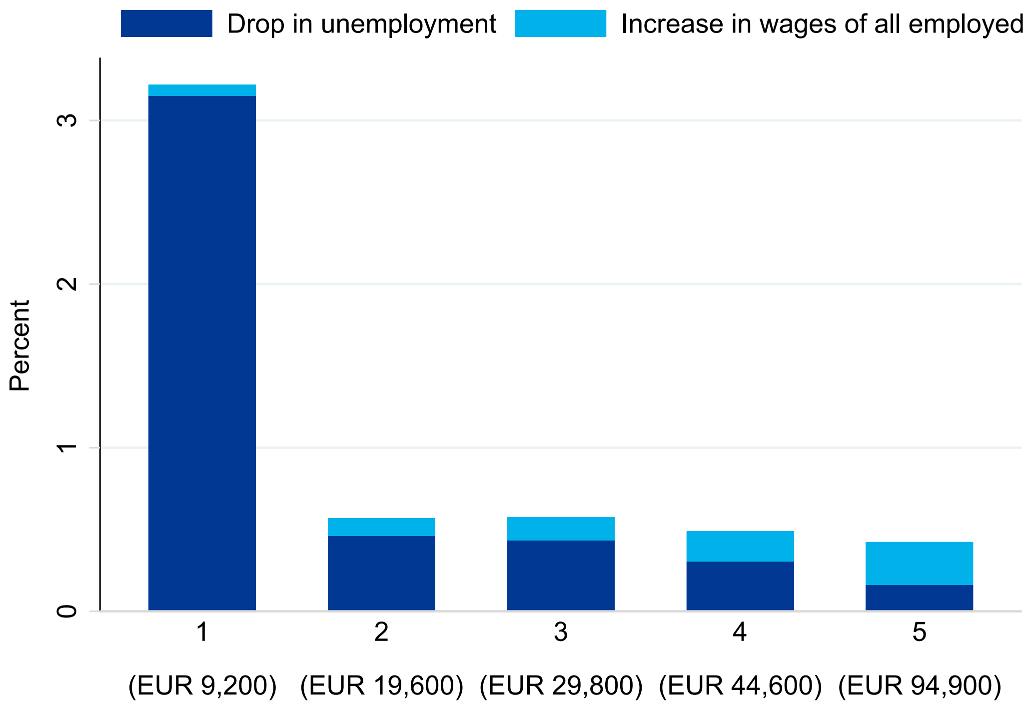

Source: Household Finance and Consumption Survey
Note: The chart shows the percentage change in mean income across income quintiles in the euro area 4 quarters after the impact of the QE shock. It also shows the decomposition of the change into the extensive margin (transition from unemployment to employment) and the intensive margin (increase in wage). The numbers in brackets show the initial levels of mean gross household income. The figure shows an aggregate of Germany, Spain, France and Italy.

Figure 7 Impulse Responses of Median Net Wealth by Country and Net Wealth Quintile

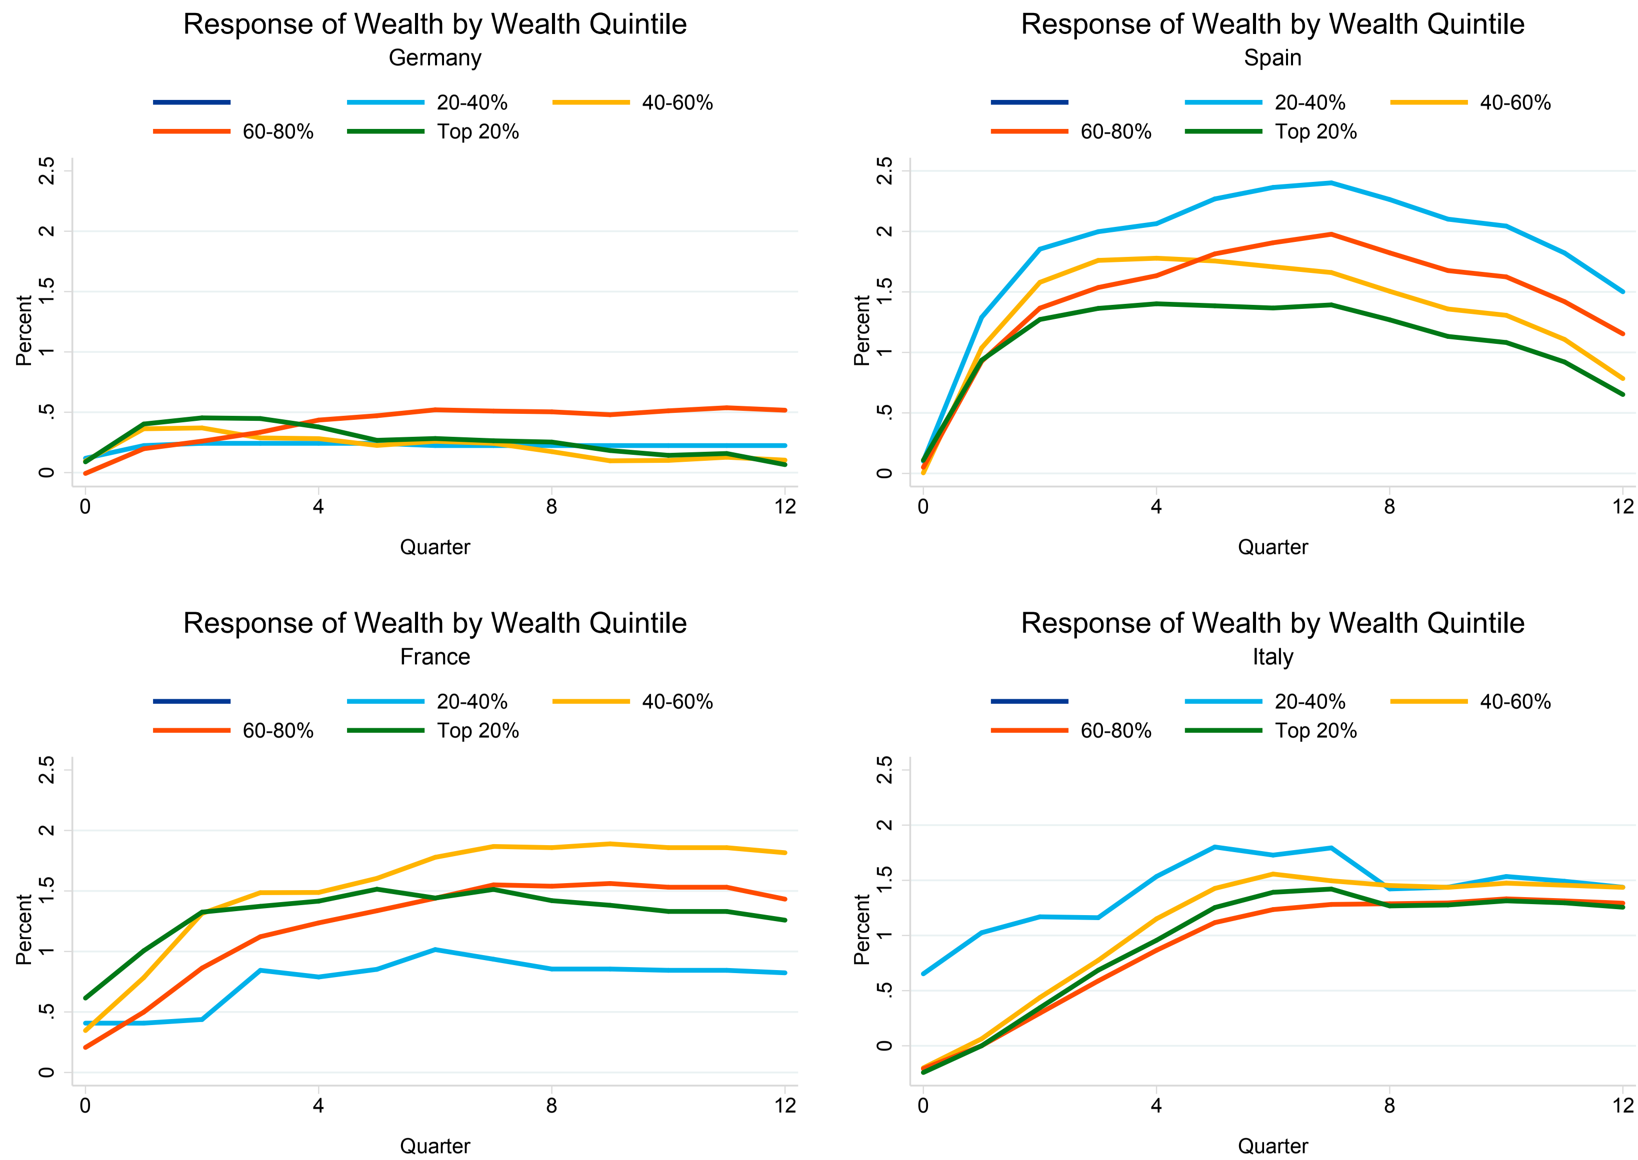

Source: Household Finance and Consumption Survey
Note: The charts show impulse responses of net wealth. The response for the bottom 20% not shown as the value of net wealth in the lowest quintile is close to EUR 0.

Figure 8 Effects of the Scenarios with Financial Income on Distribution of Income

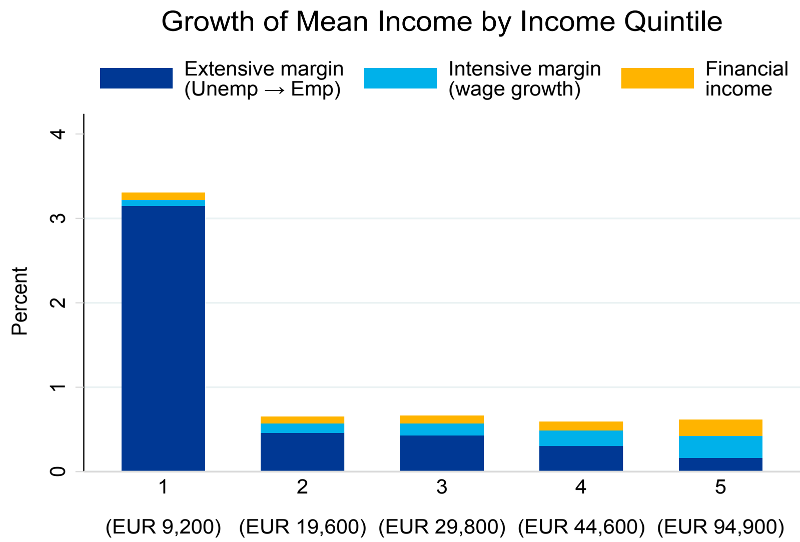

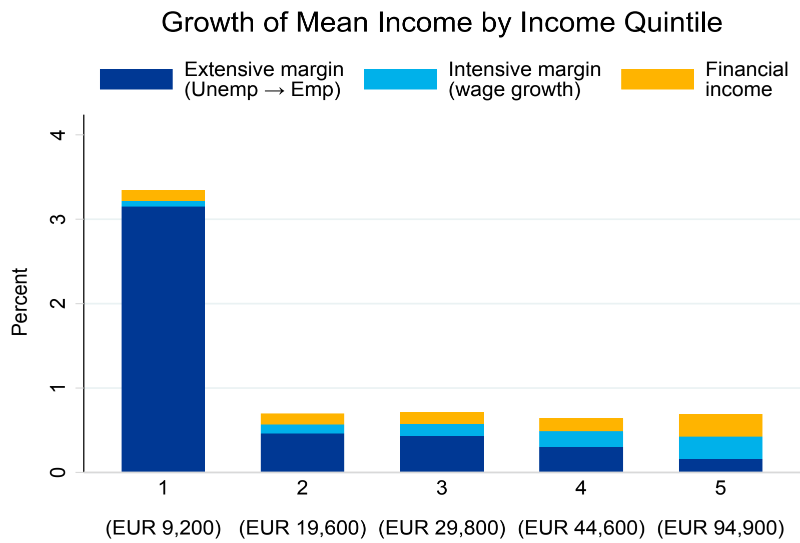

Source: Household Finance and Consumption Survey
Note: The figure shows the implications for gross income of (i) an increase of 5% in financial income (top panel) and (ii) country-specific increase in financial income (France: 6.9%, Germany: 3.6%, Italy: 19.3%, Spain: 8.3%; bottom panel). The bars show the percentage increase in mean income and its components across quintiles of gross household income. The numbers in brackets show the initial levels of mean gross household income. The figure shows an aggregate of Germany, Spain, France and Italy.

Figure 9 Effects of the Scenario with Stock Trading on the Distribution of Net Wealth

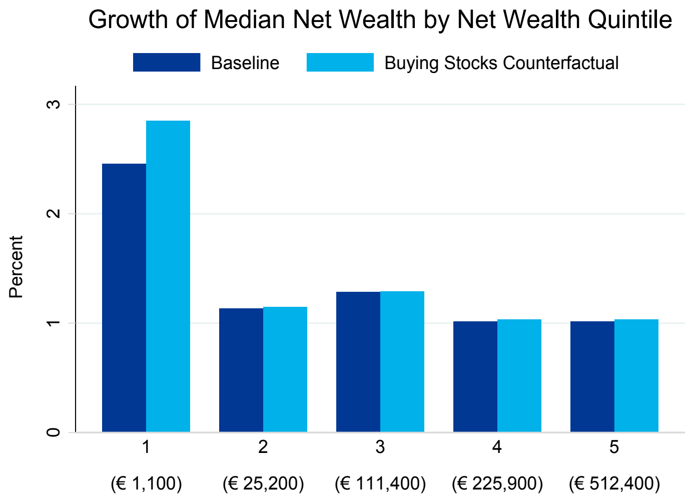

Source: Household Finance and Consumption Survey
Note: The figure compares the baseline scenario with the one in which the holding of stocks increases by 5%. The bars show the percentage increase in median net wealth across quintiles of net wealth. The numbers in brackets show the initial levels of median net wealth. The figure shows an aggregate of Germany, Spain, France and Italy.

Figure 10 Effects of the Scenario with Heterogeneity in House Price Responses on the Distribution of Net Wealth

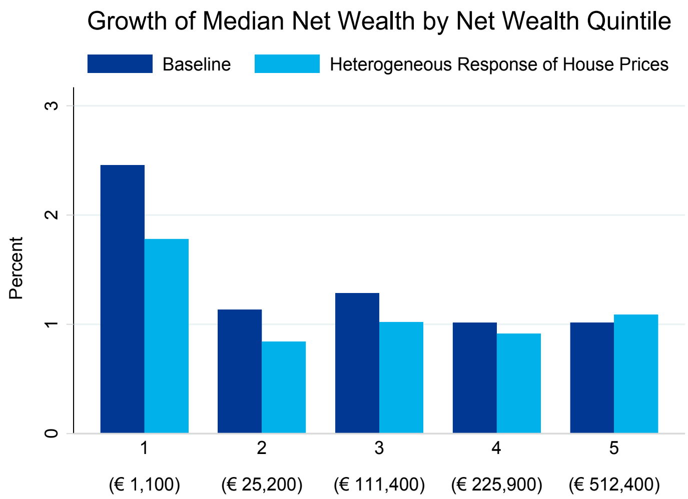

Source: Household Finance and Consumption Survey
Note: The figure compares the baseline scenario with the one in which house prices of more expensive houses (in terms of price per square meter) react more strongly to monetary policy. The bars show the percentage increase in median net wealth across quintiles of net wealth. The numbers in brackets show the initial levels of median net wealth. The figure shows an aggregate of Germany, Spain, France and Italy.

Table 1 Modeling of Responses of Wealth and Income Components at Household Level

Wealth / income component	**Modeling procedure**
Real Assets	
Household's main residence	Multiplied with response of house prices (robustness: heterogeneity in house prices)
Other real estate property	Multiplied with response of house prices (robustness: heterogeneity in house prices)
Self-employment businesses	Multiplied with response of stock prices
Financial Assets	
Shares, publicly traded	Multiplied with response of stock prices (in the baseline; robustness: some trading)
Bonds	Multiplied with response of bond prices (based on long-term rate)
Voluntary pension/whole life insurance	No adjustment
Deposits	No adjustment
Other financial assets	No adjustment
Debt	
Total liabilities (mortgage + non-mortgage debt)	No adjustment
Gross Income	
Employee income	Multiplied with response of wages (compensation per employee)
Self-employment income	Multiplied with response of wages (compensation per employee)
Income from pensions	No adjustment
Rental income from real estate property	No adjustment
Income from financial investments	No adjustment (in the baseline; robustness: grows by 5% or country-specific)
Unemployment benefits and transfers	If becomes employed, replace with wage (otherwise no adjustment)

Table 2 Effects of Quantitative Easing on Income and Wealth Inequality

	Gini Coefficient	
	Income	Net Wealth
Actual Data	43.074	68.093
Baseline Simulation	42.860	68.043
Robustness		
Effects of Financial Income (5% Response)	42.885	
Effects of Financial Income (Country-Specific Response)	42.893	
Stock Trading		68.079
Local House Prices		68.089

The table shows the Gini coefficients for gross household income and net wealth for actual data and 5 scenarios: the baseline and 4 scenarios described in section 3.2.3—2 scenarios accounting for the effects of financial income, a scenario on portfolio rebalancing of stocks (stock trading) and a scenario with heterogeneity in responses of house price to quantitative easing. The scenarios report the Gini coefficients 4 quarters after the impact of the quantitative easing shock.

Table 3 Macroeconomic Database and Identification Assumptions

Variable	Transformation	Source	QE shock	MP shock
1 Germany: GDP	log-levels	Eurostat	+	0
2 Germany: GDP deflator	log-levels	Eurostat		0
3 Germany: Unemployment rate	levels	Eurostat		0
4 Germany: House prices	log-levels	Eurostat		0
5 Germany: Compensation per employee	log-levels	Eurostat		0
6 France: GDP	log-levels	Eurostat	+	0
7 France: GDP deflator	log-levels	Eurostat		0
8 France: Unemployment rate	levels	Eurostat		0
9 France: House prices	log-levels	Eurostat		0
10 France: Compensation per employee	log-levels	Eurostat		0
11 Italy: GDP	log-levels	Eurostat	+	0
12 Italy: GDP deflator	log-levels	Eurostat		0
13 Italy: Unemployment rate	levels	Eurostat		0
14 Italy: House prices	log-levels	Eurostat		0
15 Italy: Compensation per employee	log-levels	Eurostat		0
16 Spain: GDP	log-levels	Eurostat	+	0
17 Spain: GDP deflator	log-levels	Eurostat		0
18 Spain: Unemployment rate	levels	Eurostat		0
19 Spain: House prices	log-levels	Eurostat		0
20 Spain: Compensation per employee	log-levels	Eurostat		0
21 Euro area Short-term interest rates	levels	AWM database (STN)	0	−
22 Euro area Long-term interest rates	levels	AWM database (LTN)	−	
23 Euro area Stock prices	log-levels	ECB SDW		
24 US GDP	log-levels	FRED		0
25 US Short-term interest rates	levels	FRED		0

Note: The database of the Area-Wide Model is available at `https://eabcn.org/page/area-wide-model`.