EUROPEAN CENTRAL BANK
EUROSYSTEM

Discussion Paper Series

Michele Lenza, Oreste Tristani

Heterogeneity in macroeconomics

No 27

Discussion papers

Discussion papers are research-based papers on policy relevant topics, offering a broader and more balanced perspective. While being partly based on original research, they place the analysis in the wider context of the literature on the topic. They also consider explicitly the policy perspective, with a view to develop a number of key policy messages. Their format offers the advantage that alternative analyses and perspectives can be combined, including theoretical and empirical work. The selection and distribution of discussion papers are subject to the approval of the Director General of the Directorate General Research.

Abstract

How large are the distributional effects of monetary policy in the euro area? Does heterogeneity matter for monetary policy? We answer these questions based on the results of research projects conducted at the ECB under the aegis of a dedicated research task force. A monetary policy easing causes a temporary reduction in consumption inequality; this is the case for both conventional and unconventional monetary policy. Accounting for heterogeneity at the household level provides a new perspective on the policy transmission mechanism, but it does not fundamentally alter our views of aggregate dynamics. Heterogeneity across firms appears to provide useful information to explain and forecast aggregate variables.

Keywords: Household Heterogeneity, Heterogeneity across Firms, Monetary Policy

JEL Codes: E21, E32, E43, E52, E58.

Non-Technical Summary

The experience of the past two decades, marked initially by deflationary risks and subsequently by the inflation surge, has posed new questions for monetary policy and economic modelling. Some of these questions have to do with the relevance of heterogeneity across agents (households, but also financial and non-financial firms) for the setting of monetary policy and to account for economic dynamics. Examples include: the risk of adverse effects of a prolonged period of negative interest rates, which could encourage retirement savings rather than stimulate consumption; the possible consumption boom resulting from the unwinding of the post-COVID19 excess savings, depending on the marginal propensity to consume of the holders of excess savings; and the hypothetically regressive effects of Quantitative Easing, if it proved to benefit only bond holders.

Partly as a consequence of this experience, we have witnessed the emergence of a new paradigm of macroeconomic models with heterogeneous agents. Until recently, this research had focused primarily, though not exclusively, on households and it had mostly analysed evidence from the U.S. and, due to the availability of high-quality data, from Nordic countries.

To provide new results for the euro area, related to both households and firms, a Research Task Force (RTF) on Heterogeneity was created at the ECB in 2022. This paper describes its main findings. The RTF studied the aggregate and distributional implications of household and firm heterogeneity and their implications for monetary policy. The mandate was condensed in two main questions: i) does monetary policy (or other economic shocks) produce heterogeneous effects across different agents? and ii) Is heterogeneity of economic agents relevant for monetary policy and to account for aggregate fluctuations?

The results of the RTF unambiguously point to relevant distributional implications of (conventional and unconventional) monetary policy and other economic shocks. These findings are consistent with the type of transmission channels featured in models with heterogeneous agents. While distributional implications are not the primary objective of central banks, they provide useful input for the assessment of the possible side-effects of different monetary policy instruments.

Regarding the relevance of heterogeneity for aggregate fluctuations, research on microfounded models with heterogeneous households, the so-called HANK models, including work in the RTF, has shown that they provide a more plausible description of the monetary policy transmission mechanism. More specifically, HANK models recognise that different households will be differently affected by monetary policy moves, depending on their sources of income and level of wealth. Research in the RTF has also identified cases in which shocks in HANK models produce different aggregate implications, compared to models abstracting from heterogeneity (or with a representative agent, a.k.a. RANK). Intuitively, this happens when shocks are large enough to change the

distribution of wealth and/or the models have strong HANK-specific amplification mechanisms which magnify the effects of shocks. In "normal times", the aggregate effects of macroeconomic shocks in HANK models are more similar to those obtained in RANK models.

Research in the RTF has shown that heterogeneity across firms may have the potential to change our assessment of aggregate dynamics also "in normal times". This calls for more work and data on individual firm dynamics.

1 Introduction

Monetary policy instruments are directed to the whole economy and cannot be selectively applied to particular segments of the population. This feature explains why monetary policy theory has traditionally adopted an aggregate approach, studying how central bank actions produce effects on inflation, national income, or the unemployment rate, but not on the incomes or unemployment rates of selected groups of households. The aggregate approach has been challenged in the past decade. The prolonged period of ultra-low inflation triggered a reflection on the effectiveness, and side effects, of monetary policy. Was the low interest rates period encouraging saving in some segments of the population, rather than stimulating consumption? Do Quantitative Easing (QE) policies mostly benefit financial asset holders, who are unlikely to liquidate their positions and cause an increase in aggregate demand? Such questions have led many economists to go beyond the analysis of the aggregate effects of monetary policy. Thanks to the contemporaneous development of efficient tools to solve models accounting for heterogeneous economic agents and to the increased availability of microeconomic data, this has led to an explosion of research on the distributional effects of monetary policy.

One of the key features of this literature is to provide a richer description of the monetary policy transmission mechanism to households. Starting with Kaplan et al. (2018) (henceforth KMV), the new literature underlines that only certain households, including those with savings in a bank account and those with a variable-rate mortgage, will be *directly* affected by changes in policy interest rates. Other households without savings will only react to the monetary policy decision *indirectly*, for example if they find a job after an unemployment spell, in the case of an easing decision. The distinction between direct and indirect transmission mechanisms is not only new, but also intuitively plausible. It makes frameworks accounting for household heterogeneity more compelling than the traditional, aggregate approach.

Since direct and indirect transmission mechanisms affect heterogeneous households to different extents, monetary policy does not, in general, generate a uniform impact on the entire population. Hence, monetary policy can no longer be considered as a pure stabilisation policy, because it will always produce distributional effects. Starting from this premise, this paper asks two main questions. First, how large are these distributional effects in the euro area. Second, do the distributional effects matter for the aggregate transmission of monetary policy. The paper provides an answer by summarising the results of a number of research projects conducted at the ECB under the aegis of a research task force exploring the implications of heterogeneity in macroeconomics.

Households differ both in terms of economic characteristics, such as the level and composition of their income and wealth, and of social characteristics, including education level and age. The

approach in many task force papers is to summarise this vast heterogeneity into a few dimensions that have been argued to be relevant for macroeconomics. A particularly important dimension is households' availability of liquid assets to sustain income shocks. Results show that euro area households with low liquid wealth balances adjust their consumption the most in response to a monetary policy easing. Their reaction is mostly caused by indirect effects, i.e. the general equilibrium effect of monetary policy on labour incomes. By contrast, households with larger liquidity balances respond less, and mostly through the direct impact of the policy easing on their financial returns.

By and large, these conclusions hold true for both conventional and unconventional monetary policy: the indirect effects on households with lower liquid balances is stronger both after a cut in policy interest rates, and after a reduction in long-term bond rates. Since these households have lower consumption levels on average, it follows that a monetary policy easing reduces consumption inequality. Additional evidence suggests that wage inequality falls too. From a quantitative perspective, these effects are contained and of a temporary nature.

These conclusions are relevant for the assessment of the possible side effects of central bank policies – or, in legal jargon, for a "proportionality analysis". Given that monetary policy operates through different channels, one may fear that it will benefit disproportionately, or exclusively, specific groups in the population. For example, one may be concerned that large-scale asset purchases only benefit bond holders. The results of the task force suggest that this is not the case. Expansionary monetary policy, both conventional and unconventional, produces stronger effects on more disadvantaged group in the population – those with fewer liquid balances and therefore a reduced ability to face adverse shocks.

Regarding our second question, existing estimates of the amplitude and the persistence of the response of aggregate consumption to monetary policy do not systematically change once household heterogeneity is taken into account. However, the richer set of transmission channels affecting heterogeneous households implies that, in specific contingencies, specific channels may gain importance and lead to different overall effects. Two examples appear to be relevant in the recent euro area experience.

The first one has to do with large shocks causing a sizable redistribution of wealth. A recent example is the inflation surge. Empirical work documents that this episode produced heterogeneous welfare effects across euro area households. All households lost purchasing power because of the stickiness of their nominal wage income. However, while younger and poorer households, if indebted, gained from the reduction in the real value of their debt, older and richer households suffered a loss in the real value of their assets. All in all, the inflation shock redistributed wealth

from older and richer households towards some poorer and younger households. Model-based results suggest that a redistribution of this sort would cause significant and persistent, expansionary effects on aggregate consumption, which may explain some of the surprising persistence of inflation. The reason is that the gains of younger and poorer households and the losses of older, wealthier households do not cancel out in the aggregate. Since the former households have a higher marginal propensity to consume, their net gains will translate into additional consumption expenditure, while rich households' losses will not turn into significantly lower demand. The net effect will be expansionary.

A second contingency that makes heterogeneity relevant are increases in macroeconomic or idiosyncratic uncertainty, which cause a rise in precautionary saving. An obvious example is that of unemployment risk during recessions. Its increase can generate an adverse feedback loop especially for income poor households, that are subject to higher unemployment risk and therefore contract consumption more. Model simulations suggest that this channel can be non-negligible in the case of cost-push shocks.

The different transmission mechanism of monetary policy with heterogeneous households also has implications for the optimal conduct of monetary policy. Even if the research task force has mostly abstracted from these implications, the academic literature suggests that household heterogeneity reinforces the case for a medium term orientation in monetary policy. For example, a medium-term policy horizon in reaction to adverse cost-push shocks can prevent large fluctuations in employment risk, that are particularly harmful for the poorer groups in the population.

While the main focus of this paper is on households, heterogeneity has many dimensions. Both financial and non-financial institutions also differ in terms of many features, going from size to the ease at which they have access to external financing. Some work in the task force provides an initial empirical assessments of the consequences of heterogeneity across firms. Preliminary results suggest that accounting for the distribution of firms' revenues in a reduced form model does change outcomes for macroeconomic variables. A possible theoretical channel justifying an impact of cross-firm heterogeneity on aggregate investment is that it makes productivity endogenous to monetary policy. More work focusing on heterogeneity across firms appears desirable.

We conclude the paper with a brief account of research on the implications of heterogeneity for financial stability. Since financial stability measures can more easily be targeted to individual borrowers, or to different banks, their distributional effects are a feature, not a by-product. Research shows that accounting for the interaction between heterogeneous firms and banks is important for a realistic calibration of banks' capital requirements. Rich models of the mortgage market can be used to identify those households for which borrower-based measures are most likely to be

binding.

The work surveyed in this paper focuses on the two-way relationship between heterogeneity among economic agents and monetary policy, taking as given the degree of heterogeneity observed in the data. The paper does not review some of the research projects developed at the ECB to understand the root causes of certain observed dimensions of heterogeneity. For example, Le Blanc et al. (2025) studies the determinants of the different home-ownership rates observed across euro area countries and attributes them to heterogeneous cross-country beliefs on the future evolution of house prices and to differences in the quality of the rental market. Other papers, including Savoia (2023) and Pesce and Zhang (2024), uncover new dimensions of heterogeneity in households' marginal propensity to consume.

The paper is organised as follows. Section 2 reviews the results of research articles analysing the distributional effects of monetary policy in the euro area. The typical approach in the literature is to study the effects of monetary policy *shocks*, which allow the researcher to identify the transmission mechanism of monetary policy. This is also the approach adopted in some of the papers in the task force, focusing both on conventional and unconventional monetary policy. Other papers have also analysed the impact of different shocks, because their effects are useful to identify relevant dimensions of heterogeneity in the economy. Section 3 focuses on the implications of heterogeneity for the aggregate transmission of monetary policy. This section identifies a few contingencies which make heterogeneity across household especially relevant – for example large shocks, such as the inflation surge, which cause a non-negligible redistribution of wealth. It also underlines that, over "normal" cyclical fluctuations, household heterogeneity does not fundamentally change aggregate consumption dynamics. Section 4 draws some implications for central bank policies. It briefly summarises some optimal policy results from the academic literature and then moves on to review work under the task force focusing on heterogeneity across euro area firms. This work suggests that accounting for cross-firm heterogeneity may have a systematic impact on aggregate economic dynamics, for example because changes in the composition of aggregate output can affect aggregate productivity in the economy. The section reviews also some work on the implications of heterogeneity for financial stability. We draw concluding remarks in Section 5.

2 The distributional implications of economic shocks

This section reviews the empirical work assessing the impact of macro-economic shocks on the distribution of euro area households. Our main interest is on the consumption distribution, both because this is a proxy of economic welfare at the household level, and because our main focus in

section 3 will be on the impact of the distribution on aggregate consumption. Nevertheless, we also discuss in this section the impact of macro shocks on the distribution of income and wealth, because they are building blocks to understand the effects on consumption. Among macro shocks, the main focus is on monetary policy shocks, both conventional and unconventional, but we also review results regarding the effects of other shocks producing material distributional effects, such as the inflation surge of 2021-22.

2.1 The distributional effects of conventional policy

Some early evidence on the distributional effects of conventional monetary policy in the euro area is provided in Slacalek et al. (2020). The paper quantifies the various transmission mechanisms of monetary policy to household consumption in the four largest countries. The methodology accounts for many observed sources of heterogeneity across households, going from their marginal propensities to consume, to the size and composition of their wealth and the exposure of their incomes to aggregate fluctuations. The paper argues that a simple summary of all these dimensions of heterogeneity can be obtained by grouping households according to their holdings of liquid and illiquid wealth. More specifically, following Kaplan et al. (2014), the paper focuses on three main household groups: the non hand-to-mouth, that are typically homeowners and also own sufficient liquid funds to withstand temporary income shocks; the wealthy hand-to-mouth, that are normally also homeowners but often with a mortgage and only limited liquid funds; and the poor hand-to-mouth, that are typically renters with limited liquid assets.

Table 1: Effects of a 100-bp monetary policy easing on consumption after 1 year (percent per year)

	Poor HtM	Wealthy HtM	Non HtM	All
Direct channel	0.04	0.30	0.35	0.32
Indirect channel	0.92	1.32	0.14	0.31
Total	0.96	1.62	0.49	0.63

Legend: HtM: Hand-to-Mouth. Source: Slacalek et al. (2020)

Using a simple theoretical approach, Slacalek et al. (2020) finds that households with low liquid wealth adjust their consumption the most in response to an unexpected reduction in monetary policy interest rates. Their reaction is mostly caused by the general equilibrium effects of the

monetary easing on labour incomes. Non hand-to-mouth households respond less, and mostly through the direct impact of the policy easing on their financial returns – see Table 1.

Slacalek et al. (2020) also documents significant cross-country differences between the four largest euro area countries, depending on the home-ownership rate, the prevalence of fixed- vs adjustable-rate mortgages and the volatility of the employment rate.

In terms of information, Slacalek et al. (2020) relies on the Household Finance and Consumption Survey, a triennial survey (see for example HFCN (2023)). Given the strong link between liquid wealth and marginal propensity to consume, data on liquid wealth is a "sufficient statistic" to gauge the importance of idiosyncratic factors for the effects of monetary policy. Higher-frequency microeconomic information on the health of household finances would allow much more precise and timely inference.

Jasova et al. (2021) studies the transmission of monetary policy shocks with a special focus on its impact on wage inequality. The paper relies on a granular administrative dataset for Portugal that matches employee and employer data for individual firms over the 1999–2013 period. This rich dataset allows the authors to disentangle the evolutions of between-firm and within-firm inequality. The results show that a monetary policy easing reduces wage inequality, since labour income increases more for low-wage workers than for high-wage workers. The paper also finds that these outcomes are determined by a reduction of wage inequality between firms. Surprisingly, however, these developments are accompanied by an increase in wage inequality within firms.

To understand these effects, the paper uses loan-level information on Portuguese firms from the credit registry. The reduction of wage inequality across firms after a monetary policy easing is mainly the result of the relaxation of financial constrained for small, young, and typically low-pay firms. This allows small firms to increase the wages of their workers more, compared to larger and older firms. Within each firm, however, wage increases are higher for more skilled workers, or workers with a college education. Therefore, easier monetary policy increases within-firm wage inequality because it amplifies the skill premium. All in all, financial frictions and the credit channel appear to be important determinants of the heterogeneous impact of monetary policy across workers.

den Haan et al. (2024) studies the effects of expansionary monetary policy on savings for retirement. The paper points out that expectations of persistently low policy rates, hence also persistently low returns on pension savings, imply lower consumption upon retirement. If households have sufficiently low elasticity of intertemporal substitution, they will try to smooth their consumption by increasing their current savings. In theory, this outcome may be observed for

households at a specific stage of their life cycle.

All these results underscore the view that monetary policy affects consumption through heterogeneous channels and that different households respond more or less to policy impulses, depending on which channels they are most exposed to. They are also broadly in line with results for other countries and obtained through different methodologies – see for example the survey in McKay and Wolf (2023b).

2.2 The distributional effects of large scale asset purchases

Compared to the number of existing studies on the distributional effects of interest rate policy, very little is known on the impact of QE on different households. Lenza and Slacalek (2024) is one of the few papers tackling this question with regards to the euro area.

The channels of transmission of QE are less well understood than those of conventional monetary policy. It is therefore more difficult to make a decomposition between direct and indirect effects, as in the case of Table 1. Nevertheless, Table 2 shows the total estimated effects of a reduction in 10-year government bond yields by 30 basis points on households' consumption. Rather than focusing on the hand-to-mouth characteristic, Lenza and Slacalek (2024) differentiates across households on the basis of their holdings of liquid assets. Since hand-to-mouth households are defined as households with limited liquid assets, they are likely to be found in the bottom two quintiles of Table 2.

The key result in Table 2 is that QE stimulates more strongly the consumption of households with few liquid assets. The main reason is that QE leads to a larger increase in employment rates and incomes for this group of households, who additionally have high marginal propensities to consume. Households also benefit from an increase in wealth following the QE intervention, but the implications for consumption are quantitatively smaller.

Table 1: Effects of a 30-bp drop in 10-year yields on consumption after 1 year
(percent per year)

	Lowest	20-40	40-60	60-80	Highest
Total	0.27	0.20	0.09	0.07	0.08

Note: column headers refer to quintiles of liquid wealth. Source: Lenza and Slacalek (2024)

Lenza and Slacalek (2024) also provides results on the impact of QE on the distribution of income and wealth. QE reduces income inequality because incomes increase the most at the bottom of the distribution. In turn, this is the case because of a reduction in the unemployment rate, which is particularly strong in the lowest segment of the income distribution. QE also increases financial and rental incomes, which benefit more households at the top of the income distribution. However, this effect is quantitatively small. As for the impact of QE on wealth inequality, the paper estimates essentially no impact on the Gini coefficient.

2.3 The distributional effects of other shocks

Monetary policy shocks are especially interesting, because they are a means to isolate the monetary policy transmission mechanism. However, other economic developments may also produce differential effects across households, and possibly warrant a different central bank response than if they affected all households in a uniform manner. In this section, we review the available results on the heterogeneous effects of other structural shocks and of other salient economic developments.

2.3.1 Inflation shocks

A particularly interesting recent development is the so-called "inflation surge". After three decades of price stability, inflation increased markedly and unexpectedly in the years following the COVID-19 pandemic, up to a peak of 10.6% in 2022. Which households bear the largest costs of unexpected inflation?

It is important to emphasise that the 2022-23 surge was not a "pure" inflation shock, i.e. a parallel increase in the prices of all goods. The surge was to some extent an adjustment in relative prices, because energy and food prices were the main drivers of aggregate dynamics. As a result, the distribution of welfare costs associated with the 2022-23 episode may not be representative of the costs that could be experienced in any other inflation increase. Nevertheless, many large inflationary episodes have a cost-push component, so the surge can provide useful lessons.

Pallotti et al. (2024) relies on a tractable dynamic framework to account for the various mechanisms through which household welfare can be affected by surprise inflation. Three main mechanisms appear to have been active in the euro area during the inflation surge. First, all nominal balances were permanently affected by the shock. This mechanism acted symmetrically between net borrowers and net savers. The former benefited from the real devaluation of their outstanding debt; the latter suffered a fall in the real value of, for example, their bank deposits. The second

mechanism is related to the slow adjustment of wages, rents and asset prices to the new price level. This caused a transitory fall in many real incomes, as well as in the the real price of various assets including housing. Third, due to the change in relative prices, households were differentially affected depending on the composition of their consumption baskets.

All in all, the welfare effects of inflation were starkly heterogeneous. Pensioners were the most adversely affected, with intertemporal losses up to 30% or more of annual income and caused, in all cases, by the devaluation of net nominal wealth and, in the case of low- or middle-income pensioners, by their spending a large share of their income on energy and food. By contrast, a few younger households emerged as net winners. Overall, around one quarter of euro area households gained from the inflation shock.

These results raise immediate questions as to how monetary policy ought to respond to the inflation surge. If some households gain from the inflation surge, can the monetary policy response be tailored solely to aggregate inflation and an aggregate measure of slack? Should differences in consumption baskets be taken into account? We discuss some preliminary answers to these questions in section 4.

2.3.2 Other macro shocks

Shocks to a frictional banking sector can have heterogeneous effects across households.[1] This is especially true for large shocks, or financial crises. Mendicino et al. (2024) documents that declines in U.S. bank equity returns increase consumption heterogeneity, since the consumption of low-income households decreases by roughly twice as much as that of the average household. Declines in bank returns are also associated with an increase in credit spreads and a reduction in investment, labour earnings, and asset prices. The paper demonstrates that the evidence is consistent with a model in which households at the bottom of the income distribution suffer from both losses in labour income and the increase in spreads on consumer credit, while high-income consumers can take advantage of temporarily low asset prices.

Another example of macroeconomic shocks are increases in macroeconomic uncertainty, such as the Great financial crisis or the COVID-19 pandemic. Oh and Rogantini Picco (2024) studies their effects on the U.S. population, documenting two facts: first, a heterogeneous reaction of consumption, which is cut in a significant way only by lower-income households; second, a strong reaction of aggregate labor market variables, namely an increase in unemployment and a fall in

[1] Based on Danish household data, Faccini et al. (2024) shows that indebted households' consumption reacts to increases in spreads on consumer credit. In a HANK model accounting for this feature, countercyclical variations in credit spreads induce consumption inequality.

the job-finding rate. The paper argues that these effects can be explained by a double precautionary motive. For households, and especially lower-income households, there is an increase in precautionary saving due to the rise in unemployment risk. For firms, there is a fall in vacancies, due to fears of a drop in demand. These two concomitant mechanisms can reinforce each other and generate an adverse feedback loop for income poorer households, that contract consumption the most.

Abbritti and Consolo (2024) documents the importance of skill heterogeneity and of the segmentation of labour markets across skills in the euro area. The paper shows that accounting for these features in a theoretical model increases inequality in employment, wages and consumption in reaction to technology shocks.

3 Heterogeneity and the transmission of monetary policy

Section 2 documented the distributional implications of macroeconomic shocks in the euro area. Households are differently exposed to shocks depending on the source of their income, the assets they own, and their indebtedness level. In this section we move on to discuss how household heterogeneity affects the dynamics of aggregate consumption. This question can simply not be studied in models based on the "representative household" approach, such as the New Keynesian (or NK) model, where distributional effects are nonexistent by assumption.

Recent computational advances have however made it possible to build new Keynesian models allowing for realistic features of household heterogeneity, or Heterogeneous Agents New Keynesian models (or HANK, see KMV). The key feature of HANK models is market incompleteness: on the one hand, it is acknowledged that households face very large idiosyncratic shocks, such as the risk of becoming unemployed, or being subject to other income losses or gains; on the other hand, households face borrowing constraints preventing them from borrowing against their future incomes, should negative risks materialise. After experiencing an adverse shock, households will therefore become hand-to-mouth, i.e. their consumption will be limited by their current resources.

The key question we ask in this section is: do HANK models change the aggregate transmission of monetary policy, compared to the new Keynesian benchmark? Overall, our answer is that this is only the case in specific situations. Market incompleteness and idiosyncratic risks do represent additional sources of economic fluctuations at the individual household level, but they are often not reflected in aggregate consumption dynamics. However, there are contingencies in which idiosyncratic developments do affect the aggregate economy. We provide a few illustrative examples of such contingencies.

3.1 Aggregate dynamics and idiosyncratic shocks

For concreteness, we rely on a framework which directly relates a HANK-type model to the traditional NK model. Various authors have discussed this relationship in the literature – see for example KMV or Bilbiie (2020). Here we adopt the approach in Berger et al. (2023) (henceforth BBD), which applies to broad forms of HANK models. While BBD relies on a nonlinear framework, we find it more transparent to work with approximate relationships, as this is standard in the representative-agent literature. More specifically, we follow Debortoli and Galí (2024b) and approximate aggregate dynamics to first order, and idiosyncratic dynamics to second order.[2]

In the benchmark new Keynesian model, it is well known that aggregate demand (more specifically, aggregate consumption) dynamics can be characterised by two equations (see Woodford, 2003; Galí, 2015)

$$\hat{c}_t = \mathrm{E}_t \hat{c}_{t+1} - \sigma^{-1} \left(\hat{\imath}_t - \mathrm{E}_t \pi_{t+1} \right)$$

and

$$\pi_t = \kappa \left(\phi + \sigma \right) \hat{c}_t + \beta \mathrm{E}_t \pi_{t+1}$$

where $\hat{c}_t$, $\hat{\imath}_t$ and π_t are aggregate consumption, the policy interest rate and inflation (in log-linearised form), respectively, and σ, κ and β are positive parameters. The above representation includes no shocks, but in practical applications shocks are typically added to the model to improve the fit of the data. Such shocks are interpreted as shifters in "aggregate demand", or in overall production "costs".

Based on the linearised version of the representation of heterogeneous-agent models proposed in BBD, the appendix shows that aggregate economic dynamics in a HANK model can be written as

$$\hat{c}_t = \mathrm{E}_t \hat{c}_{t+1} - \sigma^{-1} \left(\hat{\imath}_t - \mathrm{E}_t \pi_{t+1} \right) - \delta^*_{s,t} \tag{1}$$

and

$$\pi_t = \kappa^* \left(\phi + \sigma \right) \hat{c}_t + Q^*_s \mathrm{E}_t \pi_{t+1} - \kappa^* \phi \hat{\omega}_t \tag{2}$$

where $\hat{\delta}^*_{s,t}$ and $\hat{\omega}_t$ are idiosyncratic sources of fluctuation, or "wedges", which resemble shocks in the standard model, the s in $\hat{\delta}^*_{s,t}$ denotes the idiosyncratic state at time t, and the stars indicate that the equations hold for the household with the highest valuation either of a short-term bond paying interest $\hat{\imath}_t$, or of equity in the economy's firms.

[2] When HANK models are solved numerically, it is also customary to approximate their aggregate dynamics to first order – see Auclert et al. (2021). Idiosyncratic dynamics are normally solved non-linearly.

In turn, the two wedges can be expressed as

$$\delta^*_{s,t} \equiv (\mathrm{E}_t \hat{c}_{t+1} - \hat{c}_t) - (\mathrm{E}_t \hat{c}_{s,t+1} - \hat{c}_{s,t}) + \frac{1}{2}\sigma \mathrm{Var}_t \hat{c}_{s,t+1} \tag{3}$$

and

$$\hat{\omega}_t = \frac{\sigma}{\phi}\hat{c}_t + \hat{\xi}_t \tag{4}$$

where $\hat{\xi}_t$ is the second order approximation of $\Xi_t = \int_s e_{s,t}^{1+\frac{1}{\phi}} c_{s,t}^{-\frac{\sigma}{\phi}} \mathrm{d}\Phi_s$ and $e_{s,t}$ are idiosyncratic productivity shocks.

The wedge $\hat{\delta}^*_{s,t}$ looks like a demand shock in the aggregate representation. It captures the patience of the marginal asset holder, that is either the agent with the highest precautionary saving motive, or the agent with the highest expected future idiosyncratic income and who is keen to smooth her consumption over time. Equation (3) shows that patience $\hat{\delta}^*_{s,t}$ can increase in the conditional variance of idiosyncratic consumption. When it does, precautionary saving increases and consumption falls. The reduction will be unexpected – it will look like a residual – to someone looking only at aggregate data and ignoring idiosyncratic factors. Hence, a possible source of demand shocks in a representative-agent economy can be due to dynamics across heterogeneous households – see also Nakajima (2005).

For example, suppose that there are only two idiosyncratic states, a hand-to-mouth state H in which households consume their entire income, and a saver state S in which households can hold financial assets. The asset holder will be the saver and equation (3) will boil down to

$$\hat{\delta}_{S,t} = -\sigma\left(\hat{c}_t - \mathrm{E}_t \hat{c}_{t+1}\right) + \sigma\left[\hat{c}_{S,t} - p_{S,S}\mathrm{E}_t \hat{c}_{S,t+1} - (1 - p_{S,S})\mathrm{E}_t \hat{c}_{H,t+1}\right]$$

where $p_{S,S}$ is the probability to remain in the saver state and we assumed for simplicity that steady state consumption levels $c_H = c_S$. Assuming that consumption is temporarily higher in the saver state, $\hat{c}_{S,t} > \hat{c}_{H,t}$, a higher probability to become hand-to-mouth $p_{S,H} = 1 - p_{S,S}$, i.e. a fall in $p_{S,S}$, will lead to more precautionary saving by the saver and a fall in aggregate consumption. Conversely, an increase in expected future consumption in either state will induce the saver to smooth her consumption over time, i.e. to increase current consumption. Aggregate consumption will also increase.

Similarly, the wedge $\hat{\omega}_t$ looks like a cost-push shock. The term $\hat{\xi}_t$ in equation (4) shows that fluctuations in $\hat{\omega}_t$ are due to the interaction of firms' labour demand with labour supplied by workers with different idiosyncratic productivity levels. If, for example, the average of individual productivity levels defined in Ξ_t falls, firms will need to hire more labour efficiency units and $\hat{\omega}_t$

will increase. From equation (2), this change will look as a cost-push residual to someone looking only at aggregate data and ignoring idiosyncratic shocks.

Equations (1) and (2) provide some lessons regarding the impact of heterogeneity on aggregate consumption.

The first lesson is that heterogeneity will affect first-order macroeconomic dynamics if idiosyncratic developments are of an order of magnitude comparable to aggregate dynamics. Intuitively, this will only be true in specific contingencies. BBD estimates the variable $\hat{\delta}^*_{s,t}$ for the U.S. using microeconomic data and argues that one such contingency occurred when the policy interest rate reached its effective lower bound. At that point, idiosyncratic factors became the key determinants of the depth and persistence of the Great Recession. In normal times, however, idiosyncratic factors will cancel each other out in the aggregate. Thus, heterogeneity will normally not alter the evolution of the macro economy. BBD estimates that $\hat{\delta}^*_{s,t}$ normally accounts for only 7% of the fluctuations of aggregate output. As demonstrated by Debortoli and Galí (2024a), the impact of heterogeneity on aggregate fluctuations is also dependent on the specification of monetary policy in the models. When it follows a Taylor rule, which provides a realistic description of monetary policy in normal times, the role of heterogeneity for aggregate fluctuations is contained.

The second lesson is that a key contribution of heterogeneous-agent models is to provide a more fundamental explanation for some sources of economic dynamics. If one is interested in understanding *why* $\hat{\delta}^*_{s,t}$ and $\hat{\omega}_t$ are subject to change, then one must investigate the determinants of such residual. When idiosyncratic change, such as an increase in idiosyncratic risk, is plausibly underway in the economy, one must look at microeconomic evidence to validate this conjecture. If, however, one is willing to accept that $\hat{\delta}^*_{s,t}$ and $\hat{\omega}_t$ are simply exogenous and unpredictable residuals, then one will not be interested in looking at microeconomic information. For example, VAR-based empirical analyses solely relying on aggregate variables may normally uncover correct impulse responses of macroeconomic aggregates to economy-wide shocks.

3.2 Contingencies in which heterogeneity matters for aggregate dynamics

Given the importance of precautionary saving in HANK models, increases in risk are the obvious contingencies in which heterogeneity may be important for aggregate dynamics.

This possibility is investigated in Consolo and Hänsel (2024), which extends a medium-scale HANK model to incorporate search-and-matching frictions in unemployment. Such frictions generate countercyclical unemployment risk, so that recessions may be exacerbated by an increase

in precautionary saving causing a larger drop in consumption. In a model calibrated to the euro area, the paper confirms that this amplification effect is present in the data and finds that its size depends on other dimensions of heterogeneity. More specifically, if firms' profits increase during a recession, which can be the case after mark-up shocks, the contractionary effects of the recession are larger, not only because of an increase in precautionary saving by employed households fearing unemployment, but also because firms' owners will tend to save most of their larger profits. Consolo and Hänsel (2024) argues that this mechanism may have been important during the post-pandemic wave of disruptions in global supply chains and increases in energy prices.

Another contingency in which household heterogeneity matters has to do with the distribution of wealth. This is irrelevant for aggregate dynamics in a model with a large mass of identical ("representative") households, where redistributing wealth from household A to household B implies that A will consume less and B will consume more, but average consumption will remain unchanged. This is not always the case in heterogeneous-agent models. If B is initially poorer than A and subject to a borrowing constraint, then B will have a higher marginal propensity to consume than A. As a results, the increase in B's consumption following the wealth redistribution will be larger than the reduction in A's consumption. Aggregate consumption will rise.

A practical example of this conceptual experiment was observed following the recent inflation surge. As already documented above, in the euro area younger and poorer households suffered for the loss in purchasing power of their wage income, but, if indebted, they also gained from the reduction in the real value of their debt Pallotti et al., 2024. By contrast, older and richer households exposed to nominally denominated assets suffered a loss in real wealth. Pallotti (2024) studies the implications of this wealth redistribution in a simple HANK model and shows that poorer households with long-term mortgages benefit substantially from the permanent reduction in the real value of their debt. Once their nominal wages adjust to inflation, they are able to consume more, while rich households' losses do not turn into significantly lower demand. As a result, the redistribution triggers non-negligible expansionary effects on aggregate consumption and, in turn, in inflation. These effects are highly persistent and can last years, because the distribution of wealth moves very slowly in the data.

This type of redistribution is of little relevance for typical economic cycles, in which inflation fluctuations are very limited. However, it becomes sizable in case of large inflation shocks, such as the one experienced with the inflation surge. Pallotti (2024) provides some U.S.-based evidence corroborating these theoretical conclusions. Between 2019 and 2023 consumption growth has been on average higher for mortgage holders than for non-mortgage holders. Moreover, international data on credit and debit card spending suggests that countries with higher household indebtedness tended to record higher spending growth following the inflation surge.

Another illustration of the importance of the distribution of wealth is the forecasting of developments in excess savings, such as those accumulated during the lockdowns and restrictions to activities implemented after the outbreak of the COVID19 pandemic. The slow depletion of excess savings observed in the data can only be predicted by a HANK-type model (Bardoczy et al. (2024); see also Auclert et al. (2024b) for the underlying notion of *intertemporal* marginal propensity to consume).

All in all, in normal circumstances idiosyncratic developments tend to be of limited relevance for macroeconomic dynamics. While inequality can amplify cyclical fluctuations through cyclical precautionary saving (Bilbiie et al. (2024)), models with heterogeneity tell a similar story about the business cycle as representative-agent models (Bayer et al. (2024)). However, idiosyncratic developments become important in specific contingencies.[3] This is likely to occur in the case of large shocks, specifically shocks that affect the distribution of wealth, or when idiosyncratic shocks interact with other features of heterogeneity, such as the distribution of income. These considerations suggest that it would be important to account for non-linearities also in HANK models. Kase et al. (2024) develops a methodology to solve and estimate HANK models with nonlinear constraints and aggregate uncertainty.

4 Heterogeneity and central bank policies

In this section, we discuss how paying more attention to heterogeneity can better inform the preparation of monetary policy decisions. We conclude the section with some more tentative implications for financial stability policy.

4.1 Household heterogeneity: Implications for monetary policy

The work of the task force has almost exclusively focused on positive analyses, but the academic literature has also begun exploring optimal policy implications.[4] While a comprehensive review of this literature is outside the scope of this paper – for a recent perspectives see, for example, Auclert et al. (2024a) – we highlight here a few lessons emerged so far.

A general lesson is that it is impossible to make an exclusive assignment of policies to either stabilisation or redistribution goals, because every stabilisation policy has redistributive implica-

[3] As already mentioned above, the zero lower bound is another such contingency identified by BBD – see also Alves and Violante (2024).

[4] Amongst others see Bilbiie (2024), McKay and Wolf (2023a), Bhandari et al. (2021), Nuño and Thomas (2022).

tions (and viceversa) – Violante (2021). This general principle also applies to monetary policy. A corollary is that optimal monetary policy is closely dependent on the behaviour of fiscal policy – Le Grand et al. (2024). If fiscal policy had a sufficiently rich menu of instruments, there would be no redistributive role left for monetary policy. In practice, however, fiscal policy instruments are limited. Do the redistributive effects documented in section 2 warrant a change in the conduct of monetary policy?

Intuitively, one would conclude that no major modifications are necessary to a flexible inflation targeting approach. On the one hand, results in the task force and in the economic literature have focused on the cyclical relationship between household heterogeneity and monetary policy. The absence of a long-run relationship between monetary policy and inequality is taken as given. Hence, price stability should remain a desirable long-run goal for central banks. On the other hand, regarding the short-run, we have documented that a monetary policy easing tends to reduce inequality. Since monetary policy tends to be relatively more accommodative during recessions, it will 'automatically' reduce inequality when this is more desirable, i.e. when the economy is slowing down (and viceversa when the economy is strong). Taking also into account that the cyclical effects of monetary policy on inequality are contained, there is probably no need for major modifications of the standard, flexible inflation targeting approach.

By and large, formal academic analyses confirm this intuition. Heterogeneity does affect monetary policy's short-run trade-offs, but mostly in the direction of reducing the implausibly large weight on inflation stabilisation prescribed by simple theoretical models. This is the case when consumption heterogeneity is cyclical – Bilbiie (2024) – or when consumption *risk* is countercyclical and it is therefore desirable to mitigate its increase during recessions – Acharya et al. (2023). Equivalently, one can conclude heterogeneity reinforces the case for a medium-term orientation of monetary policy: after an adverse cost-push shock, the inflation target should not be reached in the shortest possible, but more gradually, in order to mitigate the costs of a too aggressive approach. Traditionally the costs are interpreted in terms of excessive unemployment. The HANK literature points out that the can also be interpreted in terms of mitigating households inequality.

A specific form of household heterogeneity which was highlighted during the inflation surge and is consequential for policy has to do with differences in consumption baskets. Bobasu et al. (2024) proposes a model to capture this feature and studies their implications for monetary policy by comparing outcomes across different monetary policy rules. Bobasu et al. (2024) emphasises that a more aggressive interest rate hike in response to inflation has both advantages and disadvantages in terms of macroeconomic adjustment. On the one hand, it is detrimental for aggregate demand, because it amplifies the reduction in consumption. On the other hand, it has the advantage of speeding up the recovery, since it allows households to rebuild their wealth more quickly thanks

to higher asset returns. However, a more aggressive response magnifies the unequal effects of the shock, because low-wealth households have little or no savings and can therefore not benefit from the speedier recovery following a more aggressive interest rate hike.[5]

4.2 Heterogeneity across firms

While the literature on household heterogeneity is growing fast, there has been less work focusing on heterogeneity across firms.

Lenza and Savoia (2024) asks whether the performance of a model of macroeconomic dynamics is improved, when information on individual firms' revenues is added to the model over and above standard information from aggregate variables. The results suggest that microeconomic information on firms is valuable: first, it leads to the conclusion that a typical business cycle is less persistent than normally thought (based on purely aggregate variables); second, it leads to a superior forecasting performance for most macroeconomic variables. These results are in contrast with those obtained for U.S. consumption with comparable methods (Chang and Schorfeide, 2024). After taking into account possible time variation in the unemployment rate, microeconomic information on consumption does not appear to change the transmission of monetary policy.

Chițu et al. (2023) also studies the heterogeneous impact of monetary policy shocks and global risk shocks on U.S. firms, but its focus is on funding costs. The paper finds no clear evidence of heterogeneity after a monetary policy tightening, suggesting that it is transmitted quite evenly to firms' funding costs. By contrast, global risk shocks determine a disproportionate increase in the funding costs of the least profitable firms.

More empirical work focusing on heterogeneity across firms appears desirable. At the same time, theoretical work provides possible mechanisms to explain the link between heterogeneity and aggregate economic dynamics. For example, González et al. (2024) studies whether firms' investment responds heterogeneously to changes in monetary policy. The paper proposes a model in which firms face a borrowing constraint, which is binding for the most productive firms and therefore implies that the allocation of capital is suboptimal in the economy. Based on Spanish data, the paper finds that firms with high marginal return on capital increase their investment more than low-return firms after an expansionary monetary policy shock. This evidence is consistent with the model and it suggests that a policy easing leads more productive firms to hold a larger share of

[5] The optimal policy response to an increase in the price of necessity goods, such as energy and food as in the case of the inflation surge, is analysed extensively in Olivi et al. (2024). The paper concludes that the response should be initially quite accommodative, so that the policy interest rate should be held relatively low, and more restrictive later on.

the aggregate capital stock. Consequently, heterogeneity across firms produces aggregate effects, because it makes productivity endogenous to monetary policy. Since more productive firms invest more after a monetary policy easing, the productivity in the economy increases.

These results hold only for temporary changes in policy interest rates. By contrast, a highly persistent, secular reduction in all interest rates, such as the one observed in recent decades across the globe, could lead to a misallocation of capital and a reduction in productivity and output – Asriyan et al. (2024).

4.3 Heterogeneity: Some implications for financial stability

Since financial stability measures are often targeted to individual borrowers, or to a specific type of financial institution, their heterogeneous effects are a feature, not a by-product. It is therefore essential to take some form of heterogeneity into account for the design of prudential policy.

A key question for financial stability is the determination of optimal capital requirements. Intuitively, this decision should take into account the risk of banks defaulting because of bad loan performance, i.e. as a consequence of the default of their borrowers. Mendicino et al. (2025) analyses a model with a tractable degree of heterogeneity, both across banks and across firms. At any point in time the model allows for two groups of firms: those that default, because they are hit by bad shocks, and those that survive. Since banks are assumed to specialise in lending to particular group of firms, they will also be subject to bankruptcy risk, depending on their capitalisation and on the performance of their loan portfolio. Abnormally high default rates among firms can therefore lead to an increase in bank failures, which constrains bank lending and leads to adverse implications for the real economy. The model calls for significantly higher optimal capital requirements than in models that ignore the impact of borrowers' default on banks' default.

Castellanos et al. (2024) studies instead the impact of borrower-specific macroprudential instruments, such as loan-to-value and loan-to-income ratios. The paper does not aim to determine optimal values of these instruments, but it provides a realistic description of the rental and housing markets. Rich dimensions of household heterogeneity are taken into account, including age, income, and wealth. Households can also be renters, homeowners, or landlords and available homes are of different sizes. A key feature that is relevant for macroprudential policy is that, upon purchasing a home, households can borrow through long-term mortgages for which macroprudential constraints must hold at origination.

The framework is tailored to the recent experience of Ireland, where minimum loan-to-income and

loan-to-value constraints were introduced in 2015. It allows the author to identify the households for which these policy measures are likely to be binding. These include mostly renters, young households and those in the bottom and middle of the income distribution.

5 Concluding remarks

We have reviewed internal ECB research studying the distributional effects of monetary policy in the euro area and drawing implications of agent-heterogeneity for central bank policies. Based on the existing work, our conclusions for monetary policy are that household heterogeneity changes our understanding of the policy transmission mechanisms, but it does not, except for specific contingencies, significantly alter the dynamics of aggregate consumption. Initial results accounting for heterogeneity across firms suggest that it may prove to be more consequential for our perceptions of key features of the economic cycle.

These conclusions are tentative, because the literature on heterogeneous agents is fast-growing. Differences across households and across firms can be observed along many dimensions, but their aggregate implications remain imperfectly understood. Macroeconomic models accounting for heterogeneity are a useful addition to the toolbox for monetary policy analysis. These models can provide especially useful insights at the time of deeper cyclical fluctuations, for example following large shocks or during episodes in which the policy interest rate reaches the effective lower bound.

References

Abbritti, M. and Consolo, A. (2024). Labour market skills, endogenous productivity and business cycles. *European Economic Review*, 170:104873.

Acharya, S., Challe, E., and Dogra, K. (2023). Optimal monetary policy according to hank. *Americal Economic Review*, 113:1741–82.

Alves, F. and Violante, G. L. (2024). Some like it hot: Monetary policy under okun's hypothesis. mimeo.

Asriyan, V., Laeven, L., Martin, A., Van der Ghote, A., and Vanasco, V. (2024). Falling interest rates and credit reallocation: Lessons from general equilibrium. *Review of Economic Studies*.

Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, page 2375–2408.

Auclert, A., Rognlie, M., and Straub, L. (2024a). Fiscal and monetary policy with heterogeneous agents. mimeo.

Auclert, A., Rognlie, M., and Straub, L. (2024b). The intertemporal keynesian cross. mimeo.

Bardoczy, B., Sim, J., and Tischbirek, A. (2024). The macroeconomic effects of excess savings. Finance and economics discussion series 2024-062. washington: Board of governors of the federal reserve system, https://doi.org/10.17016/feds.2024.062.

Bayer, C., Born, B., and Luetticke, R. (2024). Shocks, frictions, and inequality in us business cycles. *American Economic Review*, 114(5):1211–47.

Berger, D., Bocola, L., and Dovis, A. (2023). Imperfect risk sharing and the business cycle. *Quarterly Journal of Economics*, 138:1765–1815.

Bhandari, A., Evans, D., Golosov, M., and Sargent, T. (2021). Inequality, business cycles and monetary-fiscal- policy. *Econometrica*, 89:2559–2599.

Bilbiie, F., Primiceri, G., and Tambalotti, A. (2024). Inequality and business cycles. mimeo.

Bilbiie, F. O. (2020). The new keynesian cross. *Journal of Monetary Economics*, 114:90–108.

Bilbiie, F. O. (2024). Monetary policy and heterogeneity: an analytical framework. *Review of Economic Studies*, page 1–39.

Bobasu, A., Dobrew, M., and Repele, A. (2024). Energy price shocks, monetary policy and inequality. ECB Working Paper No. 2967.

Castellanos, J., Hannon, A., and Paz-Pardo, G. (2024). The aggregate and distributional implications of credit shocks on housing and rental markets. ECB working paper No. 2977.

Chang, M. and Schorfeide, F. (2024). On the effects of monetary policy shocks on earnings and consumption heterogeneity. NBER Working Paper 32166.

Chițu, L., Grothe, M., Schulze, T., and Van Robays, I. (2023). Financial shock transmission to heterogeneous firms: the earnings-based borrowing constraint channel. ECB Working Paper 2860.

Consolo, A. and Hänsel, M. (2024). Hank faces unemployment. ECB Working Paper 2953.

Debortoli, D. and Galí, J. (2024a). Heterogeneity and aggregate fluctuations: Insights from tank models. In Eichenbaum, M. S., Leahy, J. V., and Ramey, V. A., editors, *NBER Macroeconomics Annual*, volume 39. Forthcoming.

Debortoli, D. and Galí, J. (2024b). Idiosyncratic income risk and aggregate fluctuations. *American Economic Journal: Macroeconomics*, 16:279–310.

den Haan, W., Ferrari, A., Mazelis, F., and Ristiniemi, A. (2024). The role of unemployment risk, aging, and retirement for business cycles. mimeo.

Faccini, R., Lee, S., Luetticke, R., Ravn, M., and Renkin, T. (2024). Financial frictions: Macro vs micro volatility. mimeo.

Galí, J. (2015). *Monetary Policy, Inflation and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press, Princeton.

González, B., Nuño, G., Thaler, D., and Albrizio, S. (2024). Firm heterogeneity, capital misallocation and optimal monetary policy. ECB Working Paper 2890.

HFCN (2023). Household finance and consumption survey: Results from the 2021 wave. European central bank statistics papers series no. 46.

Jasova, M., Mendicino, C., Panetti, E., Peydro, J.-L., and Supera, D. (2021). Monetary policy, labor income redistribution and the credit channel: Evidence from matched employer-employee and credit registers. CEPR Discussion Paper No. 16549.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108:697–743.

Kaplan, G., Violante, G. L., and Weidner, J. (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*, pages 77–138.

Kase, H., Melosi, L., and Rottner, M. (2024). Estimating nonlinear heterogeneous agent models with neural networks. mimeo.

Le Blanc, J., Slacalek, J., and White, M. N. (2025). Housing wealth across countries: The role of expectations, institutions and preferences. CEPR Discussion Paper 19838.

Le Grand, F., Martin-Baillon, A., and Ragot, X. (2024). Should monetary policy care about redistribution? optimal monetary and fiscal policy with heterogeneous agents. mimeo.

Lenza, M. and Savoia, E. (2024). Do we need firm data to understand macroeconomic dynamics? Sveriges Riksbank Working Paper Series 438.

Lenza, M. and Slacalek, J. (2024). How does monetary policy affect income and wealth inequality? evidence from quantitative easing in the euro area. *Journal of Applied Econometrics*, 39:746–765.

McKay, A. and Wolf, C. (2023a). Optimal policy rules in hank. mimeo.

McKay, A. and Wolf, C. K. (2023b). Monetary policy and inequality. *Journal of Economic Perspectives*, 37:121–44.

Mendicino, C., Nikolov, K., Rubio Ramirez, J., Suarez, J., and Supera, D. (2025). Twin defaults and bank capital requirements. *Journal of Finance*.

Mendicino, C., Nord, L., and Peruffo, M. (2024). Distributive effects of banking sector losses. mimeo.

Nakajima, T. (2005). A business cycle model with variable capacity utilization and demand disturbances. *European Economic Review*, 49:1331–1360.

Nuño, G. and Thomas, C. (2022). Optimal redistributive inflation. *Annals of Economics and Statistics*.

Oh, J. and Rogantini Picco, A. (2024). Macro uncertainty, unemployment risk, and consumption dynamics. *International Economic Review*.

Olivi, A., Xhani, D., and Sterk, V. (2024). Optimal monetary policy during a cost-of-living crisis. mimeo.

Pallotti, F. (2024). The fisher channel according to hank: Unexpected inflation and the missing recession. mimeo.

Pallotti, F., Paz-Pardo, G., Slacalek, J., Tristani, O., and Violante, G. L. (2024). Who bears the costs of inflation? euro area households and the 2021–2023 shock. *Journal of Monetary Economics*, 148:103671.

Pesce, S. and Zhang, L. (2024). Mortgage refinancing and the marginal propensity to consume. mimeo.

Savoia, E. (2023). The effects of labor income risk heterogeneity on the marginal propensity to consume. ECB Working Paper 2866.

Slacalek, J., Tristani, O., and Violante, G. L. (2020). Household balance sheet channels of monetary policy: A back of the envelope calculation for the euro area. *Journal of Economic Dynamics and Control*, 115:103879.

Violante, G. L. (2021). What have we learned from hank models, thus far? Proceedings of the 2021 ECB Forum on Central Banking.

Woodford, M. (2003). *Interest and Prices: Foundations of a theory of monetary policy*. Princeton University Press, Princeton.

A Appendix

A.1 Households

Temporary utility for household s at time t is

$$u\left(c_{s,t}, l_{s,t}\right) = \left(\frac{1}{1-\sigma} c_{s,t}^{1-\sigma} - \chi \frac{l_{s,t}^{1+\phi}}{1+\phi}\right)$$

where $c_{s,t}$ is consumption of a composite good and $l_{s,t}$ is labour supply.

Households can trade in a short term bond $M_{s,t}$ with nominal return I_t and in shares $\Omega_{s,t}$ with price P_t^{Ω}. The budget constraint is

$$P_t c_{s,t} + M_{s,t} + \Omega_{s,t} P_t^{\Omega} \leq V_{s,t} + l_{s,t} e_{s,t} w_t P_t + T_{s,t}$$

where P_t is the price index, $T_{s,t}$ are lump sum transfers, $e_{s,t}$ is household productivity, w_t is the real wage for efficiency unit, and $V_{s,t}$ is total nominal financial wealth at the beginning of the period, given by $V_{s,t} = I_{t-1} M_{s,t-1} + \Omega_{s,t-1}\left(P_t^{\Omega} + P_t d_t\right)$, for real dividends d_t.

Following BBD, we assume trading restrictions on equity holdings such that

$$H\left(M_{s,t}, \Omega_{s,t}\right) \geq 0$$

for $H_M \geq 0$, so that holding more bonds relaxes the trading restriction on shares.

The first order conditions include

$$\frac{1}{I_t} \geq \mathrm{E}_t \beta \left(\frac{c_{s,t+1}}{c_{s,t}}\right)^{-\sigma} \frac{1}{\Pi_{t+1}} \tag{5}$$

where the greater-or-equal sign is due to the shadow value of the trading restriction $H\left(M_{s,t}, \Omega_{s,t}\right) \geq 0$. If the restriction is slack, $H_{M_{s,t}} = 0$ and equation (5) holds with equality. The corresponding household is the one valuing the bond the most (or accepting the lowest possible return).

Assuming that there is always a household for which (5) holds with equality, we proceed as in Debortoli and Galí (2024b) and approximate the equation to first order with respect to aggregate variables and to second order with respect to household-specific variables. We obtain

$$\hat{i}_t = \mathrm{E}_t \pi_{t+1} - \sigma\left(\hat{c}_{s,t} - \mathrm{E}_t \hat{c}_{s,t+1}\right) - \frac{1}{2}\sigma^2 \mathrm{Var}_t \hat{c}_{s,t+1} \tag{6}$$

Note that this can be rewritten as

$$\frac{1}{\sigma}\left(\hat{i}_t - \mathrm{E}_t \pi_{t+1}\right) + \frac{\sigma}{2}\mathrm{Var}_t \hat{c}_{s,t+1} = \mathrm{E}_t \Delta \hat{c}_{s,t+1}$$

which corresponds to equation A2 in the appendix of Debortoli and Galí (2024b).

Equation (6) can be rewritten as

$$\hat{c}_t = \mathrm{E}_t \hat{c}_{t+1} - \sigma^{-1}\left(\hat{i}_t - \mathrm{E}_t \pi_{t+1}\right) - \delta^*_{s,t}$$

that corresponds to equation (1) in the main text, for

$$\delta^*_{s,t} \equiv \left(\mathrm{E}_t \hat{c}_{t+1} - \hat{c}_t\right) - \left(\mathrm{E}_t \hat{c}_{s,t+1} - \hat{c}_{s,t}\right) + \frac{1}{2}\sigma \mathrm{Var}_t \hat{c}_{s,t+1}$$

Note that an increase in the variance of idiosyncratic consumption leads to an increase in precautionary saving and a fall in aggregate consumption.

A.2 Firms

There is a continuum of firms of measure one. The production function for firm k is

$$Y_{k,t} = Z_t h_{k,t}$$

where $h_{k,t}$ is the demand of labour services and Z_t is aggregate productivity. Given the utility function, aggregate demand for good $Y_{k,t}$ is the standard Dixit-Stiglitz aggregator $Y_{k,t} = \left(\frac{p_{k,t}}{P_t}\right)^{-\theta} Y_t$, where $p_{k,t}$ is the price of good $Y_{k,t}$. In each period, firm k's nominal profits are

$$\Pi_{k,t} = (1-\tau)\, p_{k,t} Y_{k,t} - P_t w_t h_{k,t}$$

where τ is a sales tax.

Firms choose their price $p_{k,t}$ subject with Calvo probability $1-\alpha$. Following BBD, we assume that firms are held by the household that values dividends in a given aggregate state next period the most. If we define the corresponding marginal rate of substitution as $Q^*_{s,t,T}$, firm k's maximisation problem is

$$\max_{p_{k,t}} \mathrm{E}_t \sum_{T=t}^{\infty} \alpha^{T-t} Q^*_{s,t,T} \left[(1-\tau)\, p_{k,t} \left(\frac{p_{k,t}}{P_T}\right)^{-\theta} Y_T - P_T \frac{w_T}{Z_T} \left(\frac{p_{k,t}}{P_T}\right)^{-\theta} Y_T \right]$$

with standard first order condition

$$\frac{p_{k,t}}{P_t} = \frac{\theta}{\theta - 1} \frac{\mathrm{E}_t \sum_{T=t}^{\infty} \alpha^{T-t} Q^*_{s,t,T} w_T \left(\frac{P_t}{P_T}\right)^{-\theta-1} \frac{Y_T}{Z_T}}{\mathrm{E}_t \sum_{T=t}^{\infty} \alpha^{T-t} Q^*_{s,t,T} (1-\tau) \left(\frac{P_t}{P_T}\right)^{-\theta} Y_T}$$

Labour services are equal to the aggregate bundle of labour skills in the economy. Total labour demanded by firms is therefore

$$\int_0^1 h_{k,t} \mathrm{d}k = \int_s e_{s,t} l_{s,t} \mathrm{d}\Phi_{s,t}$$

which can be rewritten as $\frac{Y_t}{Z_t} \Delta_t = \int_s e_{s,t} l_{s,t} \mathrm{d}\Phi_{s,t}$ for $\Delta_t = \int_0^1 \left(\frac{p_{k,t}}{P_t}\right)^{-\theta} \mathrm{d}k$.

Using the intratemporal condition, wages can be written as $w_t = \chi \left(\frac{Y_t}{Z_t} \frac{\Delta_t}{\xi_t}\right)^{\phi}$ where

$$\Xi_t = \int_s e_{s,t}^{1+\frac{1}{\phi}} c_{s,t}^{-\frac{\sigma}{\phi}} \mathrm{d}\Phi_s$$

Note that, in a representative household model, Ξ_t would simplify to $\Xi_t = c_t^{-\frac{\sigma}{\phi}}$. The Phillips' curve can therefore be linearised around a zero inflation steady state as

$$\pi_t = \kappa^* \left[(\phi + \sigma)\, \hat{c}_t - (1+\phi)\, z_t \right] + Q^*_s \mathrm{E}_t \pi_{t+1} - \kappa \phi \hat{\omega}_t$$

for $\kappa^* \equiv \frac{(1-\alpha)(1-\alpha Q_s^*)}{\alpha}$ and

$$\hat{\omega}_t = \frac{\sigma}{\phi}\hat{c}_t + \hat{\xi}_t$$

where $\hat{\xi}_t$ is the second order approximation of Ξ_t, and we imposed $\hat{y}_t = \hat{c}_t$. Ignoring for simplicity the aggregate productivity shock z_t, we obtain equation (2) in the main text.

Acknowledgements
This paper summarises the work done as part of an ECB's Research Task Force set up to explore the implications of agent-heterogeneity in macroeconomic and finance. We wish to thank Agostino Consolo, Jordi Galí, Luc Laeven, Caterina Mendicino, Filippo Pallotti, Gonzalo Paz-Pardo, Jiri Slacalek and Gianluca Violante for useful comments and suggestions.
The opinions in this paper are those of the authors and do not necessarily reflect the views of the European Central Bank and the Eurosystem.

Michele Lenza
European Central Bank, Frankfurt am Main, Germany; Centre for Economic Policy Research, London, United Kingdom;
email: michele.lenza@ecb.europa.eu

Oreste Tristani
European Central Bank, Frankfurt am Main, Germany; Centre for Economic Policy Research, London, United Kingdom;
email: oreste.tristani@ecb.europa.eu

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

PDF ISBN 978-92-899-7457-8 ISSN 2811-8820 doi: 10.2866/1748926 QB-01-25-229-EN-N