Monetary Policy and Inequality[†]

Alisdair McKay
FRB Minneapolis

Christian K. Wolf
MIT & NBER

January 3, 2023

Abstract: We ask three questions about the connection between monetary policy and inequality. First, does monetary policy affect inequality? While different households respond to changes in monetary policy for different reasons, we argue that the overall consumption effects are relatively evenly distributed across households. Second, does household heterogeneity change our understanding of monetary policy transmission? A more careful account of microeconomic consumption behavior materially alters our understanding of transmission channels, but has rather limited effect on our general view of the aggregate effects of monetary policy. Third, does inequality affect the optimal conduct of monetary policy? Since monetary policy is a rather blunt distributional tool, we argue that even a central bank with an explicit distributional mandate would not deviate much from conventional policy prescriptions.

[†]Alisdair McKay is a Senior Research Economist, Federal Reserve Bank of Minneapolis, Minneapolis, Minnesota. Christian K. Wolf is Assistant Professor of Economics, Massachusetts Institute of Technology, and a Faculty Research Fellow, National Bureau of Economic Research, both in Cambridge, Massachusetts. Their email addresses are `alisdair.mckay@mpls.frb.org` and `ckwolf@mit.edu`. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. The authors thank Adrien Auclert, James Cloyne, Martin Blomhoff Holm, Pascal Paul, Joe Vavra, Johannes Wieland, and Abbie Wozniak for helpful comments.

1 Introduction

Monetary policy affects incomes, employment rates, asset prices, and interest rates. Because different households hold different financial assets, work in different sectors, and are differentially attached to the labor force, it follows that monetary policy will almost certainly have heterogeneous effects across households. The last decade has seen an explosion of empirical and theoretical research on the links between monetary policy and inequality. In this article, we take stock of this important research agenda. We organize our discussion around three main questions.

First, how unevenly distributed are the effects of monetary policy, and why? Monetary stimulus—that is, lower nominal interest rates—affects households through many different channels. Most directly, households may pay lower interest rates on their debts and are likely to earn lower returns on their savings. In addition, changes in interest rates also set into motion changes in the broader economy that indirectly affect households: jobs become easier to find, wages and prices increase, and asset prices rise. Taking any one of these channels in isolation, the effects of a monetary policy easing can look very uneven. For example, many households will not benefit at all from an increase in stock prices, simply because they do not own any stocks. Aggregating across all of these transmission channels, however, the gains appear much more evenly distributed. Roughly speaking, low-income households benefit from a tighter labor market, middle-class households benefit from lower mortgage rates, and wealthy households benefit from capital gains on assets. To reach this conclusion, we review recent empirical work on the incidence of monetary policy across households. We also provide descriptive measures of the heterogeneity in household balance sheets and their exposure to changes in monetary policy.

Second, does a more careful account of the microeconomic effects of monetary policy affect our understanding of its propagation to the macroeconomy at large? Yes and no. Theoretical analyses that dig into the microeconomic propagation of monetary policy have materially altered our views on transmission channels. For example, we have learned that an important component of the transmission of monetary policy to consumer spending is actually indirect: by changing nominal interest rates, monetary policy first of all directly increases consumer and firm demand; this increase in demand then in general equilibrium leads to higher income, which in turn leads to meaningful second-round effects on consumer spending. At the same time, however, there is rather limited scope for such micro-level studies to change our overall view on the macro effects of monetary policy. The reason is simple: we already have good

empirical evidence on the overall response of aggregate output and inflation to changes in interest rates, and any structural model of monetary policy transmission—with or without microeconomic heterogeneity—needs to be broadly consistent with that evidence. That said, the more detailed view of the micro-level effects of monetary policy that emerges from recent research does suggest new reasons to expect the economy's sensitivity to monetary policy to vary over time and with the state of the business cycle.

Third, how—if at all—should the interaction between monetary policy and inequality affect the behavior of central banks? We argue that even if the central bank's mandate includes distributional concerns, appropriate policy is unlikely to differ too much from the optimal policy of a central bank that is solely focused on macroeconomic goals like stabilizing inflation and aggregate activity. The reason is related to our answer to the first question: because the effects of monetary policy are relatively evenly distributed, the scope for achieving distributional objectives through monetary policy is likely to be rather limited.

2 The distributional effects of monetary policy

Our objective in this section is to assess how monetary policy affects consumption across different groups of households. One could in principle instead ask how monetary policy affects other measures of inequality such as income or wealth. We will do so as well, but only with the end goal of translating income and wealth into consumption. We do so because, while of course not synonymous with utility, consumption is more directly related to a household's well-being than are its wealth and income.

At any given point in time, monetary policy will of course explain very little of the overall differences in consumption across different groups of households. Therefore we discuss how consumption *changes* across household groups following a *change* in monetary policy. These consumption responses will depend—at least to some degree—on the institutions of the economy (e.g., whether mortgages are mostly fixed-rate or floating-rate). We focus primarily on the U.S. context, though many of the forces we describe are likely to be similar across countries. We will also phrase our discussion in terms of the effects of monetary stimulus; in many respects a monetary contraction would have the mirror image effects, though there are some important sources of asymmetry that we will note. We proceed in two steps. First, we review the key channels through which monetary policy will affect consumption. Second, we combine all of those channels to arrive at overall conclusions for how monetary policy will affect the consumption of heterogeneous households.

2.1 Channels of transmission

Monetary policy affects the consumption of an individual household by changing the prices, wages, interest rates, and opportunities it faces. Many of these changes are indirect in nature; for example, expansionary policy may lead to a tighter labor market, thus resulting in higher wages for workers. This section considers some of the main channels of transmission from monetary policy decisions to household consumption decisions.[1] For each channel, we will describe its strength in the aggregate as well as its heterogeneous incidence across households. To do so, we will combine empirical evidence on the aggregate effects of changes in monetary policy with data on heterogeneity in household finances. That data is taken from the 2019 Survey of Consumer Finance (Federal Reserve Board, 2019), a nationally representative survey of households that collects data on their income, assets, and liabilities.

Throughout this section we will repeatedly refer to Figure 1. Each panel of three figures (the rows of the figure) shows a different feature of household balance sheets. The left column of the figure splits households into five wealth quintiles and reports the average within each quintile. The middle and right columns do the same splitting households into income and age quintiles, respectively. We focus on these dimensions of heterogeneity across households as they interact most closely with the mechanisms studied in the recent literature.

Income. Expansionary monetary policy stimulates the aggregate economy and thereby raises labor income. These gains are unlikely to be distributed equally, as labor incomes of low-income households tend to be disproportionately exposed to the business cycle (Okun, 1973; Guvenen et al., 2014, 2017; Patterson, 2022). This pattern holds both for business cycles in general, and also for expansions and contractions that result from changes in monetary policy in particular (Andersen et al., 2021; Amberg et al., 2021; Holm et al., 2021). The inequality in earnings resulting from this income incidence channel is quite pronounced. For example, data from the recession of 1979–1983—which was arguably caused by a monetary contraction—shows that the earnings losses for low-income households were many times larger than those for high-income households (see Guvenen et al., 2014).

These heterogeneous responses of income are then further amplified when translated to changes in consumption. In particular, low-income households are more likely to be financially constrained—that is, with no savings and no access to credit. Without a buffer of

[1] A theoretical decomposition of individual consumption responses to monetary policy that includes these channels can be found in Auclert (2019).

savings or credit, such households are then likely to have a stronger response of consumption to a change in income. To illustrate the empirical relevance of this point we in Panel (a) of Figure 1 report the fraction of households that are financially constrained. To construct the figure, we follow Kaplan et al. (2014) and classify households as financially constrained if their liquid assets amount to less than two weeks' worth of income. According to this definition, it is possible even for high net worth households to be financially constrained if their assets are held in illiquid forms (such as real estate). Overall we classify 31 per cent of households as being constrained in this fashion, with households of lower net worth and with lower income much more likely to be constrained.

Revaluation of nominal contracts. Expansionary monetary policy raises the general level of prices by temporarily raising the rate of inflation. In most borrowing and lending arrangements, the contracting parties agree to a repayment that is set in nominal terms, and so a surprise increase in the price level will reduce the real value of the repayment. Doepke & Schneider (2006) document the heterogeneity in household exposure to such surprise inflation. A typical middle-class household has substantial nominal debts in the forms of mortgages, auto loans, credit card debts, and student loans. For a typical household, their nominal assets (e.g. bank deposits and bond holdings) are smaller then their nominal debts, and so it follows that a surprise increase in inflation will lower the real value of their debts by more than it will lower the real value of their assets. This is especially true of young households—a group that tends to have large mortgage balances. Old, rich households, on the other hand, tend to have more nominal assets than nominal liabilities and so the net worth of these households declines after an unexpected increase in the price level.

How large are these effects? It is widely believed that inflation is actually fairly insensitive to short-run changes in monetary policy (for example, see Mavroeidis et al., 2014). On the other hand, since most central banks nowadays try to stabilize rates of *inflation*, it follows that the implied changes in the *price level* are permanent, with even small changes in the inflation rate accumulating over time to have a substantial effect on the overall price level. In particular, as we describe further below, a change in monetary policy that raises GDP by one percentage point for two years approximately leads to a 0.5 percentage point permanent increase in the price level. For a household with a mortgage balance that is twice their annual income (about the average ratio for young homeowners), this channel thus represents a decrease in the real value of their liabilities equal to one percent of their annual income.

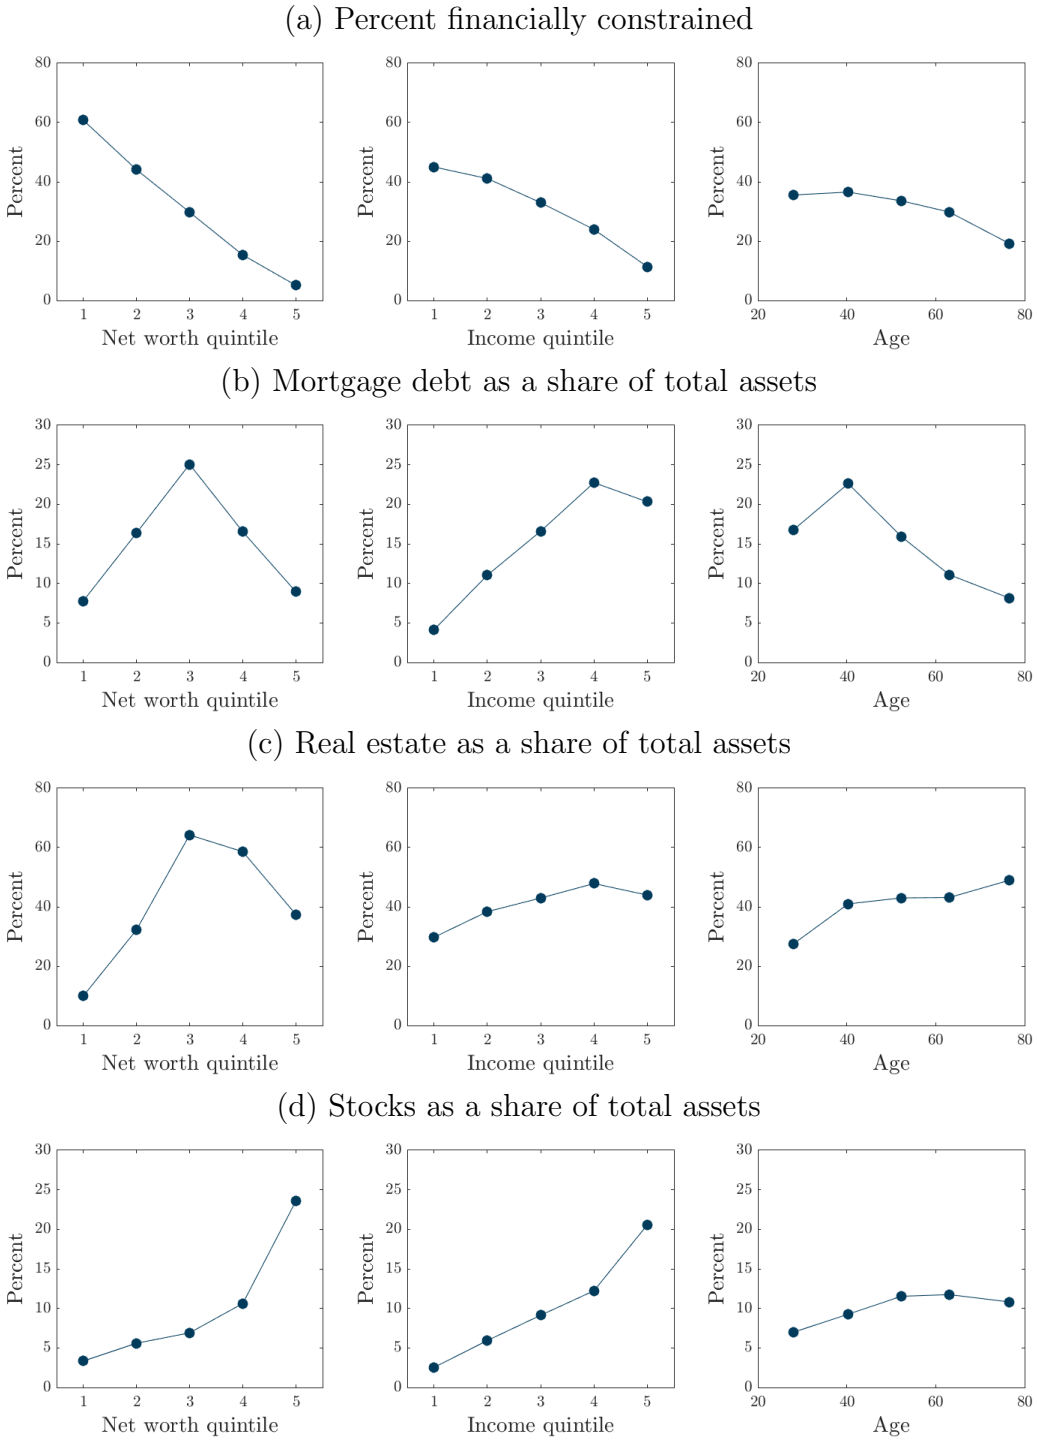

Figure 1: Summary measures of household balance sheets from the Survey of Consumer Finances (2019). Each figure divides the sample into quintiles by net worth, income, or age and then plots the mean level within that quintile. Panel (a) shows the fraction of households that have liquid assets of less than two weeks' income. Panel (b) shows the mean ratio of mortgage debt to total assets. Panel (c) shows the share of real estate in household asset holdings. Panel (d) shows the share of stocks in households assets.

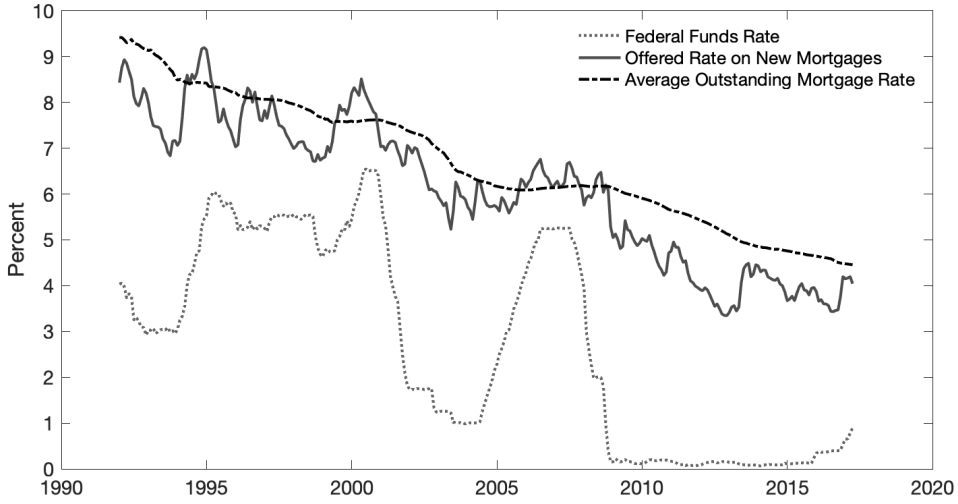

Figure 2: Monetary policy and mortgage rates. Offered rate is the average rate currently offered by mortgage lenders (source: Freddie Mac). Outstanding rate is the average rate currently paid by homeowners as computed in Berger et al. (2021). The federal funds rate is an overnight interest rate and in particular the target rate for monetary policy implementation (source: Federal Reserve).

Mortgages. A homeowner with a mortgage not only benefits from the change in the price level, but they may also benefit from a reduction in the nominal interest rate they pay. Mortgage interest rates are tightly linked to monetary policy: a one percentage point change in the federal funds rate typically translates to roughly a 0.5 percentage point change in the 30-year mortgage rate (for a review of evidence, see Wong (2021) and references therein). The design of mortgage contracts varies across countries. In the U.S., most mortgages have fixed nominal interest rates, but the borrower is free to repay the loan at the time of their choosing. This option to repay the loan creates an asymmetry: when nominal interest rates fall, homeowners can benefit by refinancing their loans at lower interest rates; if rates rise, homeowners can simply keep their original loans. Figure 2 plots the average interest rate paid on outstanding mortgages (as computed by Berger et al., 2021) along with (i) the rate currently offered on new 30-year mortgages as well as (ii) the federal funds rate. As they are long-term rates, offered mortgage rates are less volatile than the federal funds rate. Moreover, as most borrowers have fixed-rate contracts and refinancing is gradual, the rates they actually pay are even smoother than the offered rates. Finally, due to the asymmetry mentioned above, average mortgage rates tend to co-move more closely with the federal funds rate following rate cuts than rate hikes.

Panel (b) of Figure 1 shows what kind of households will tend to gain from lower interest

rates through this mortgage channel. In particular, the figure shows the ratio of mortgage debt to household assets. We see that mortgage debt is particularly important for households in the middle of the wealth distribution. Moreover, the ratio of mortgage debt to assets tends to decline with age, so young households are more exposed to changes in mortgage rates than are older households. Returning to our example of a young household with a mortgage balance of twice their income, a percentage point reduction in their mortgage rate implies a two percent increase in disposable income for the life of the loan, which is typically between five and ten years.

Of course, whenever households pay lower interest rates on their debts, somebody else is receiving less interest income. In the U.S. mortgage market, the lenders are often the owners of mortgage backed securities—financial contracts that entitle the owner to receive the principal and interest payments on mortgages. It is generally difficult to say how this lost interest income will feed back to households, as mortgage backed securities are typically held by financial institutions, governments, and foreign investors rather than outright by households.[2] As the loss of interest income is directly borne by financial institutions, governments, and foreign investors, it is rather unlikely to have a substantial effect on household consumption decisions, at least in the short run.

Asset prices. Expansionary monetary policy tends to increase the values of long-lived assets such as stocks and real estate. As real interest rates decline, the present discounted value of future cash flows increases, thus leading to a revaluation of assets and liabilities. At first glance, this channel may appear to be a key channel of monetary policy-induced redistribution: asset-holders get wealthier when rates fall. However, lower interest rates also mean lower expected returns on these assets—a force pushing in the opposite direction. Which of these effects dominates depends on the horizon at which the asset-holder plans to consume.[3] Fagereng et al. (2022) provide a simple and intuitive way for thinking about the redistributive effects of changes in asset prices: those who plan to sell the asset benefit when its price increases, and those who plan to buy the asset are harmed. Additional important effects relate to the role of assets as collateral and as buffers against changes in income: through these channels, higher asset prices may lead households to spend more even if they

[2]Data from 2010 shows that the U.S. government and foreign investors owned nearly 50% of outstanding agency mortgage backed securities (Tracy & Wright, 2012).

[3]More precisely, what matters is how a household's plan to consume at various dates lines up with the existing claims to cash flows across those dates. Households with front-loaded consumption and backloaded cash flows benefit from lower rates (Auclert, 2019; Greenwald et al., 2021; Fagereng et al., 2022).

do not plan to buy or sell the asset. In what follows we elaborate on the connection between changes in wealth and changes in consumption for two important asset classes: stocks and housing.

Stock prices are highly sensitive to changes in interest rates (see Bernanke & Kuttner, 2005; Bauer & Swanson, 2022). As discussed above, the effect of stock price changes on household consumption is likely to depend on the future savings and consumption plans of the household. Empirical evidence overall suggests that higher stock market wealth does translate to an immediate (if moderate) increase in consumer spending, with an extra dollar of stock wealth increasing consumption by 3 cents (Chodorow-Reich et al., 2021). Across households, the ownership of stocks is highly concentrated, with wealthy households holding the vast majority. Panel (d) of Figure 1 reveals that wealthy households also devote a much larger share of their portfolios to stocks.

House prices tend to increase following expansionary monetary policy (Iacoviello, 2005). Moreover, housing is particularly important as a share of household balance sheets for the middle class (e.g., see in Panel (c) of Figure 1).[4] Many homeowners expect to remain in their houses for many years; since they then do not plan to either buy or sell the asset it may seem that they are unaffected by a change in house prices. Intuitively, following a monetary easing, those households now own a more valuable house, but they now also want to live in a more expensive house. However, houses often serve as valuable sources of liquidity for households, either as collateral for loans or through the reassurance that they can rely on their home equity as a financial backstop in the future (e.g., see Berger et al., 2018). Empirical evidence suggests that homeowners as a group overall do increase their consumption when their homes appreciate–a "housing wealth effect." In particular, recent estimates show that an additional dollar of housing wealth leads to an increase in consumption of between 3 cents (Guren et al., 2021) and 7 cents (Mian et al., 2013).

Intertemporal substitution. Finally, monetary policy can induce households to substitute consumption across time. Intuitively, the real interest rate is the price of consumption today relative to consumption in the future, and a decline in this relative price should result in households increasing consumption today and reducing consumption in the future. How strongly households respond to these incentives depends on the types of consumption we

[4]Here we study asset price changes and thus we focus on homeowners rather than renters. Renters that do not plan to purchase a home are not directly affected by these asset price changes, though they could be affected indirectly if rents adjust to reflect these changes in home prices.

are considering. For non-durable goods and service consumption, empirical evidence (see Best et al., 2020) suggests that this intertemporal substitution effect is quite small; it is also homogeneous across households, thus limiting its redistributive effects.

2.2 Assessing consumption effects through direct measurement

We have seen that monetary policy affects household balance sheets through several distinct margins, with each channel likely to benefit different groups of households. Ultimately, we are interested in how the sum of these changes translates to household consumption. In the rest of this section we describe two broad approaches to answering this question: a direct approach (this subsection) and an indirect approach (the next one).

A direct measurement approach is perhaps the most natural way to proceed: that is, measure consumption at the household level and then ask how the distribution of consumption responds to changes in monetary policy. The key challenge with this approach is that high-quality data on individual consumption is needed—the data need to allow the researcher to tease out the role of monetary policy among the many other (and actually more important) factors that also affect inequality across households.

U.S. DATA. For the U.S. economy, the best source for individual consumption data is arguably the Consumer Expenditure Survey—a nationally representative survey of households that is conducted quarterly by the Census Bureau on behalf of the Bureau of Labor Statistics. Coibion et al. (2017) use these data to create quarterly time series of statistics that summarize the consumption distribution—for example, the 10th, 50th, and 90th percentiles of the consumption distribution in each quarter. They then explore how these distributional measures evolve after a plausibly exogenous change in monetary policy. They find that expansionary monetary policy reduces consumption inequality, because the 90th percentile of the distribution moves closer to the 10th and 50th percentiles. Chang & Schorfheide (2022) also use Consumer Expenditure Survey data for the same purpose, though with a somewhat different econometric approach and a different measure of monetary policy. They instead find that expansionary policy *increases* consumption inequality, with the top-end of the distribution now moving away from the rest. Overall, one of the challenges with this direct approach—and a potential reason behind the conflicting findings in prior work—is that the contribution of monetary policy shocks is small relative to other factors that affect a cross-section of households. As a result, sampling variation can be an important obstacle in isolating the heterogenous effects of policy.

One way of circumventing these challenges is to estimate heterogeneous consumption effects of monetary policy across broader groups of households, notably across homeownership status. Cloyne et al. (2020) find that the consumption levels of homeowners with a mortgage and renters respond to changes in interest rates by similar percentage amounts, while the consumption of homeowners without a mortgage does not appear to react as strongly. Specifically, they find that a one percentage point reduction in nominal interest rates leads to about a percentage point increase in non-durable consumption for mortgagors and renters, and no statistically significant spending response for homeowners without mortgages. This broad gradient by homeownership is then likely to translate into heterogeneity in the effects of monetary policy by wealth, income, and age. First, for the U.S., the fraction of households in the Survey of Consumer Finance that has a mortgage is about 60 percent for the top three wealth quintiles and rises steadily across the income quintiles, reaching 73 percent for the highest quintile. This suggests consumption gains that are somewhat increasing in wealth and income. Second, older households are less likely to have a mortgage, so younger households should benefit by more than older households. The implied gradient in consumption responses by age is consistent with results reported in Wong (2021). She finds that the consumption response to monetary policy is concentrated among homeowners who take out a new mortgage (either to refinance an existing one or purchase a new home) after the change in monetary policy.

Evidence from other countries. Another way of circumventing the challenge of sampling variation is to use an administrative dataset that contains the entire population of households. Such data is not available for the United States. However, Holm et al. (2021) construct measures of household consumption using Norwegian tax data on the incomes and assets of households. The key insight behind this approach is that household-level consumption can be imputed reasonably well by using the assets held by a household at the start of the year, adding the income received during that year, and then subtracting the assets held at the end of the year. A further key benefit of the Norwegian data is its panel structure, allowing Holm et al. to measure the *change* in consumption at the household level following a change in monetary policy.

The results of Holm et al. suggest that expansionary monetary policy has U-shaped effects on consumption across the wealth distribution, with asset-poor and asset-rich households increasing their consumption somewhat more than households in the middle. The authors find similar patterns for disposable income, reflecting strong responses of non-financial

income at the bottom of the asset distribution and financial income at the top of the distribution. The largest changes in consumption and disposable income occur with a substantial delay after the change in policy. At these horizons, the least-wealthy and wealthiest groups increase consumption by 1.5 to 2 percentage points, while those in the middle increase consumption by about one percentage point. Overall we view the Holm et al. study as particularly informative given its use of high-quality panel data. However, caution should be used in applying the Norwegian results in a U.S. context: for example, Norwegian mortgages usually have adjustable interest rates and Norwegian households have relatively few direct holdings of stocks.

2.3 Assessing consumption effects through indirect measurement

An indirect approach offers an alternative to direct measurement: by using a combination of theory and empirical evidence, one can aggregate the various individual channels of monetary transmission reviewed in Section 2.1 into a total effect on household consumption.[5] On the whole, taking into account all channels, we will argue that the consumption changes from a monetary easing appear relatively evenly distributed in the cross-section of households.

Ingredients. Table 1 presents the ingredients that we use for our indirect calculation of household consumption responses to a monetary easing. The table lists (i) a variety of "prices" to which households are exposed—i.e., the channels of Section 2.1—and (ii) how sensitive those prices are to changes in monetary policy. To construct the values in the second column we estimate how those various prices respond to a plausibly exogenous change in interest rates induced by monetary policy, with details on identification and estimation presented in Appendix A. All estimates are in real terms and have been scaled to correspond to a monetary stimulus that leads real GDP to increase by 1% on average over the first two years following the change in policy. The first row shows that labor earnings respond slightly more than GDP, while the second row reveals that business income moves about one-for-one with GDP. Further down the table, we see that stock prices are very sensitive to monetary policy, increasing by about five times more than GDP. Lower real interest rates—including in particular lower mortgage rates—reflect both lower nominal rates as well as an increase in the price level.

[5] Auclert (2019) is an important and well-known example of the indirect approach to assessing the distributional effects of monetary policy. Slacalek et al. (2020) focus on European data and pursue a closely-related approach to the one we present here.

		Marginal prop. to consume	
	Price change	Unconstr.	Constrained
Labor earnings	1.3%	0.05	0.5
Business income	1.0%	0.05	0.5
Interest income	-1.5%	0.05	0.5
Return on stocks	4.8%	0.03	0.03
Return on housing	0.6%	0.03	0.03
Return on cash	-0.56%	0.05	0.5
Mortgage rates	-0.65%	0.8	1.0
Other interest rates	-0.88%	0.8	1.0

Table 1: Inputs for indirect calculation of consumption effects of monetary policy. Each row of the table corresponds to a channel of monetary policy effects on household consumption. The table lists how strongly the price or income associated with that channel responds to monetary policy and how strongly households spend out of that income category. The constrained column applies to households with liquid assets less than two weeks' worth of income. Changes in incomes, assets returns, and interest rates are in real terms (see Appendix A for further details).

The third and fourth columns of the table list an assumed marginal propensity to consume for each type of income—that is, the strength of the consumption change following a change in income, for each income category. For example, a value of 0.05 means that a one dollar increase in income would lead to a 5 cents increase in consumption. It is important to note that a household's marginal propensity to consume is likely to depend on its (financial) circumstances; for example, according to economic theory, a household that is financially constrained may spend strongly out of any additional income, while households with access to savings or credit are more likely to save additional income or use it to pay down debt. To capture these effects in a simple and transparent way we assume that households with few liquid assets are financially constrained and have a high marginal propensity to consume out of transitory changes in income; high liquid-wealth, unconstrained households, on the other hand, will have lower marginal propensities to consume. The third and fourth columns of the table reflect this split—one for financially constrained households and one for unconstrained households. We assume that all households have high marginal propensities to consume with respect to changes in debt service payments (based on the discussion in Di Maggio et al., 2017) but a low marginal propensity to consume with respect to house and stock price appreciation (as discussed above).

Towards total consumption responses. By combining the ingredients in Table 1 we can construct our indirect estimates of cross-sectional consumption responses to monetary policy changes. We proceed as follows. For each household in our Survey of Consumer Finance dataset, we first classify them as financially constrained or unconstrained depending on their ratio of liquid assets to income, as already discussed above for Figure 1. Based on this classification we assign them the corresponding marginal propensities to consume reported in Table 1. Then, for each channel of policy transmission corresponding to a row in the table, we calculate their exposure to that channel. For example, a household's exposure to the stock market depends on the size of their stock holdings, while their exposure to mortgage rates depends on their current mortgage balance and so on.[6] For labor income, we assume that low-income households are disproportionately exposed to the labor market, consistent with the findings in Guvenen et al. (2017). For each channel we then compute a household's change in consumption as their exposure times the estimated price change (as listed in the table) times the assumed marginal propensity to consume. Summing across the different transmission channels we arrive at a total effect on the household's consumption. We report the results in terms of a *percentage* change in consumption.[7]

Results. Our calculation gives us an estimate of how each household in the Survey of Consumer Finances would change their consumption following a hypothetical monetary stimulus. We then summarize the distribution of these consumption changes in Figure 3. The figure shows the consumption effects across different quintiles of net worth (left panel), income (top right) and age (bottom right).

The left panel shows an average consumption change of 0.8 percent among low-wealth households and an average consumption change of 1.2 percent for wealthy households. The shaded areas in this left panel decompose the total consumption effect, revealing that labor income and non-mortgage debt drive the consumption response for low-wealth households, wages and mortgages are the main factors in the middle of the wealth distribution, and stock market gains are increasingly important at high wealth levels.

The top-right panel of Figure 3 instead shows the total consumption effect across income levels for working-age households. The average consumption responses within the top four

[6]To account for fixed-rate contracts and gradual refinancing we scale the change in mortgage rates and other interest rates by a factor of 1/2.

[7]To compute a percentage change, we need a baseline level of consumption. As the Survey of Consumer Finances does not report consumption, we impute it using the relationship between consumption and income in the Panel Study of Income Dynamics (Institute for Social Research, 2019).

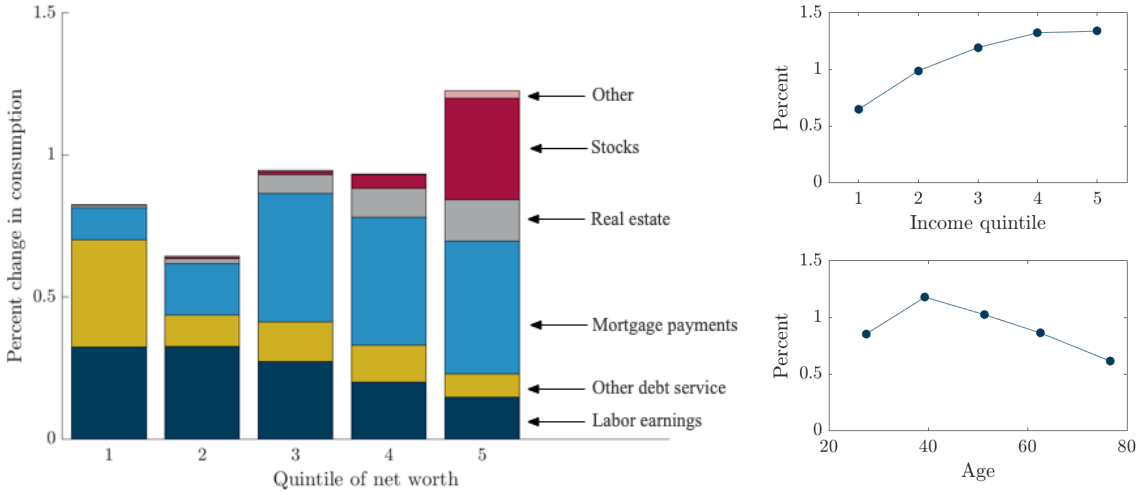

Figure 3: Indirect calculation of consumption change across quintiles of net worth, income, and age. Sample restricted to households with net worth less than $2.5 million. See Appendix A for details on the construction.

quintiles are all between 1.0 and 1.3 percent, while the lowest income quintile response is lower at 0.6 percent. Mechanically, a partial explanation for the insensitivity of consumption among the low-income group is that a substantial part (32% on average) of the income in this group is derived from social insurance and other transfers, which our analysis assumes is insensitive to monetary policy. Another potential explanation is that our analysis may actually understate the response of labor earnings in this group. Among households in the lowest income quintile, 34 per cent had zero labor income during 2018, which was the reference year for the survey data we use. As our analysis "scales up" the existing income categories, the households that start with zero labor income will by construction not gain any labor income in our calculation; in practice, however, it may well be possible that these households would in fact enter employment in an expanding economy.

Finally, the bottom-right panel of Figure 3 shows the total consumption effect across age. Consistent with the logic reviewed in Section 2.1 (which suggested that young homeowners benefit substantially), we here find that the consumption effect peaks in early middle age and declines thereafter. Another factor driving the relatively small consumption response for old households is that a substantial part of their income comes from Social Security payments, which we assume are unaffected by monetary policy.

Overall, Figure 3 suggests two main takeaways on the cross-sectional incidence of monetary policy on household consumption. On the one hand, the incidence of the *individual*

channels of monetary policy transmission to households is quite uneven. For example, mortgage payments and stocks have much stronger effects at the top of the wealth distribution, while other debt service and labor income has stronger effects at the lower end. On the other hand, once aggregated across all channels, the *overall* consumption changes are much more evenly distributed. In particular, across the various cuts of the data, *all* groups materially increase their consumption. While there are some differences across groups, we overall view them as relatively modest. In particular, after a monetary stimulus that raises total GDP by one per cent, even the least affected groups increase their consumption by a still material 0.6 per cent (vs. 1.3 per cent for the most affected).

A MORE STRUCTURAL APPROACH. While the indirect calculation above is attractive for its simplicity, it does require a large number of strong, reduced-form assumptions about household marginal propensities to consume. A more structural alternative is to use a model of household consumption decisions to infer how household consumption responds to the changes in income and prices induced by monetary policy. We do exactly that in McKay & Wolf (2022), using a general-equilibrium model in which households own a variety of long-duration assets and are unequally exposed to changes in labor earnings, in line with our discussion of transmission channels in Section 2.1. In that analysis, we also find that monetary stimulus leads to a quite evenly distributed increase in consumption across the population of households. Figure 4 summarizes our results, showing the consumption responses to monetary stimulus by net worth quintiles. The key takeaway is that, across all levels of wealth, consumption responds by nearly the same percentage amount. In this case, expansionary monetary policy roughly scales up everyone's consumption by the same amount as the aggregate, leaving each household's *share* of total consumption approximately unchanged.

3 Inequality & the aggregate effects of monetary policy

Traditionally, the transmission of monetary policy to the macroeconomy has been analyzed in models populated by a representative household that chooses aggregate consumption, savings, and labor supply (for textbook treatments see Woodford, 2003; Galí, 2015). Recent research has pursued a different approach that explicitly incorporates household heterogeneity. This research agenda starts from microeconomic modeling of the choices of individual, heterogeneous households. We then arrive at predictions for aggregate variables by summing

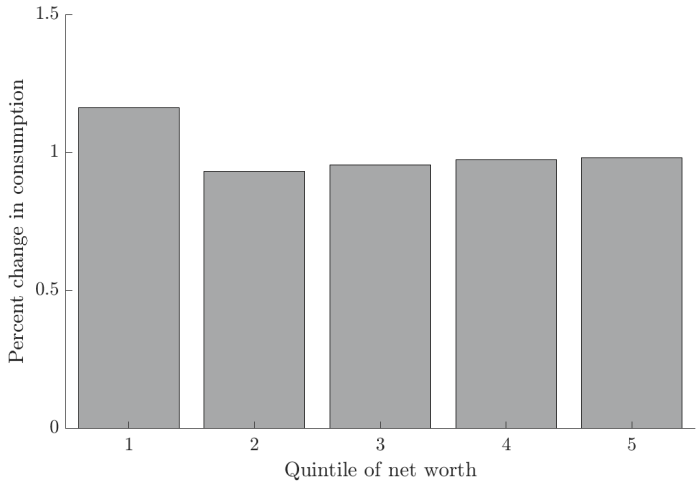

Figure 4: Consumption response to monetary stimulus at different levels of net worth from McKay and Wolf (2022).

up across these heterogeneous households. This bottom-up, heterogeneous-agent approach can account for many of the distributional channels that we discussed in the previous section.

A NEW VIEW ON THE CHANNELS OF POLICY TRANSMISSION. The heterogeneous-agent approach, with its emphasis on consumption-savings decisions at the household level, has changed our understanding of the *decomposition* of monetary policy effects into different underlying channels of transmission. Two margins of the transmission mechanism have received particular attention: (i) the role of mortgage refinancing and (ii) the consumption response to changes in income.

As we have described above, expansionary monetary policy is typically associated with a decline in mortgage interest rates. Importantly, in the U.S., homeowners often have the option to refinance their fixed-rate mortgages to take advantage of lower rates on new mortgages. As households refinance their mortgages to lower interest rates, their disposable income increases, allowing them to consume more. Propagation through such mortgage refinancing—rather than intertemporal substitution, as emphasized in traditional on macroeconomic models—thus emerges as one of the most important direct transmission channels of monetary policy to consumer spending (Beraja et al., 2019).

In general equilibrium, the extra demand induced by monetary policy then translates to tighter labor markets, decreasing the unemployment rate and increasing labor incomes. Empirical evidence on the consumption response to changes in income shows that households spend quite strongly out of such temporary income gains (see Johnson et al., 2006). Many of

the new heterogeneous-agent models of monetary policy transmission are designed to match these empirically estimated strong spending responses to changes in income. This leads to an important *indirect* channel of policy transmission: expansionary policy raises incomes and then households spend strongly out of that income, reinforcing the initial increase in demand. These indirect effects can be particularly strong due to a Keynesian multiplier logic: income increases spending, which then further increases household income, and so on (see Kaplan et al., 2018; Auclert et al., 2018; Bilbiie, 2020; Patterson, 2022). Important empirical support for these model predictions was provided in Holm et al. (2021).

Re-assessing the aggregate effects of monetary policy. The heterogeneous-agent view has changed our understanding of the precise channels through which monetary policy operates. Does this translate to a change in thinking about how monetary policy affects the macroeconomy as a whole? The answer is somewhat nuanced.

At a broad level, we would argue that our understanding of the effects of monetary policy on macro outcomes has not changed very much. This reflects the simple fact that we already have good *empirical* evidence on the average effects of monetary policy on macroeconomic outcomes (Ramey, 2016, see). Any structural model of monetary policy transmission needs to be consistent with this evidence on the total effect, and so at best microeconomic heterogeneity will affect our understanding of the decomposition of this total effect into different channels. Importantly, however, as we refine our understanding of the precise channels through which monetary policy operates, we may then also change our perspective on why policy may be more or less powerful at a given point in time depending on the state of the economy—a phenomenon known as state dependence in the effectiveness of policy. Such state dependence is difficult to identify from purely empirical analysis of time series data because it requires the researcher to estimate the effects of policy in different states of the economy (as opposed to just estimating some average effect). Structural modeling—and in particular modeling that carefully accounts for microeconomic channels of transmission—is thus the most promising avenue to learn about such state dependence.

The role of mortgage refinancing. One likely reason for state dependence in the aggregate effects of monetary policy is related to household mortgage refinancing (Berger et al., 2021; Eichenbaum et al., 2022). The incentives for households to refinance their mortgages depend on the difference between the mortgage rate offered on a new loan and their existing interest rates. If households currently have high interest rates on their mortgages, then they will be likely to refinance soon anyway, and thus any additional changes in mort-

gage rates due to monetary policy will flow through strongly to the rates households actually pay. On the other hand, if households are already paying low interest rates, they will be less likely to refinance, and so any marginal rate change related to monetary policy will have a smaller impact on the rates households are actually paying.

Beraja et al. (2019) analyze another related reason for why the strength of the mortgage refinancing channel is likely to vary over time. In order to refinance a mortgage, the homeowner must be approved for a new loan. Importantly, obtaining a new loan can be more or less difficult depending on a number of factors that are likely to vary over time. For example, some lenders require that the new loan is for no more than 80% of the home's value. In the aftermath of the Great Recession, declines in home prices left many homeowners unable to meet this requirement. These homeowners found it difficult to refinance their mortgages and were unable to take advantage of the low interest rates offered on new mortgages. In the aggregate, at times when many households are in this situation, the mortgage channel of monetary policy transmission will be muted.

Spending on durables. Time variation in household demand for durable goods is another reason for why the sensitivity of the economy to monetary policy may change over time (Berger & Vavra, 2015; Tenreyro & Thwaites, 2016). For concreteness consider a household that is contemplating the purchase of a new car. For such a household, a change in interest rates could make the difference between buying the car and not buying the car. It follows that monetary policy is likely to have large effects when many households are contemplating such purchases, as usually happens in times of economic expansion. Conversely, in a downturn, few households are contemplating any big purchases, and so monetary policy transmission may be weakened.

Intertemporal shifting of demand. Finally, in addition to raising the possibility that the power of monetary policy varies over time, the recent heterogeneous-household research agenda has also raised questions about the medium-term effects of policy. We typically think that monetary stimulus raises demand in the short run. But what about at longer horizons? Mian et al. (2021) as well as McKay & Wieland (2021) highlight forces whereby monetary stimulus raises demand in the near term but depresses it at longer horizons. The logic of these arguments is that monetary stimulus raises demand today but changes household balance sheets in ways that leave them less willing to spend in the future. This could be because they take on additional debt (Mian et al., 2021) or because they purchase durable goods (McKay & Wieland, 2021). These studies predict that changes in interest rates will

tend to be persistent because stimulus today requires continued stimulus in the future to offset the endogenous reduction in future demand.

Taking stock. Our overall conclusion is that recent research emphasizing microeconomic household heterogeneity has lead to an evolution—rather than a revolution—in our understanding of the aggregate effects of monetary policy. Compared to prior work, this research places emphasis on a different set of channels shaping the aggregate effects of policy changes. While these channels introduce some novel sources of state dependence and intertemporal shifting of demand, the existing empirical evidence on the aggregate short-run effects of monetary policy remains an important touchstone for both representative-agent as well as heterogeneous-agent models.

4 Optimal monetary policy with household inequality

So far we have discussed the interaction between monetary policy and inequality on purely positive grounds, asking whether (i) monetary policy affects the evolution of inequality, and (ii) inequality affects the propagation of monetary policy. We now turn to a normative question, asking how inequality may change our view of *optimal* monetary policy.

In the long run, real economic outcomes—including inequality across households—are largely outside the control of monetary policy. For optimal monetary policy, inflation stabilization remains the only long-run consideration. Here, we will focus instead on how monetary policy should respond to fluctuations in the economy in the short run. We will start by considering a central bank with a narrow mandate, seeking only to stabilize macroeconomic aggregates. We then consider how a central bank with a broader mandate that includes distributional concerns would act differently. Much of the intuitive discussion in this section is based on our formal analysis in McKay & Wolf (2022).

4.1 A narrow mandate

Traditionally, central banks have pursued the dual objectives of stabilizing (i) inflation as well as (ii) real aggregate activity measures (such as employment or GDP). Will household inequality affect the behavior of a central bank with these targets?

Policy problem. Figure 5 shows a stylized version of the policy problem faced by a central bank with a mandate to stabilize inflation and aggregate output. The top panel

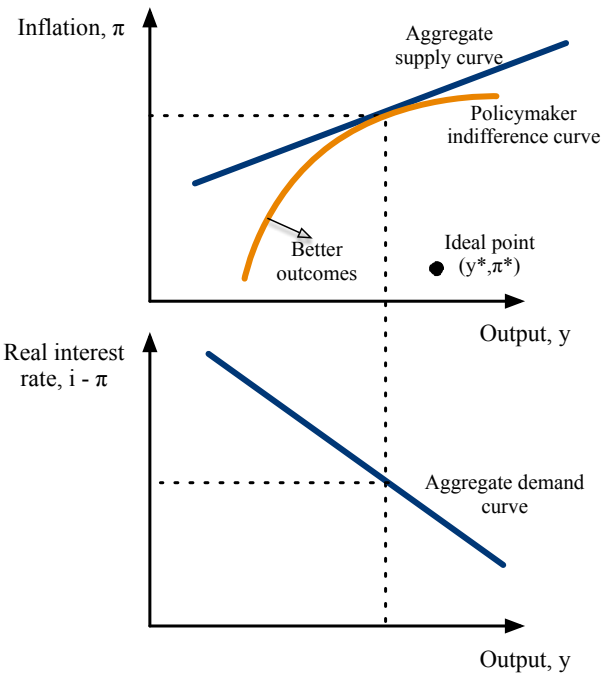

Figure 5: Optimal policy with a narrow mandate.

begins by showing policymaker preferences (in orange) and constraints (in blue). The figure features output on the horizontal axis and inflation on the vertical axis, with the output-inflation pair marked as (y^*, π^*) as the policymaker's desired outcome. The orange curve then shows an indifference curve corresponding to policymaker preferences, with better outcomes closer to the target.[8] The blue line on the other hand is the constraint set—the set of inflation-output pairs that the policymaker can in principle implement. The line corresponds to an aggregate supply curve, with its upward slope reflecting the usual logic that higher utilization of the economy's resources leads to upward pressure on costs and thus prices. Intuitively, if the economy is going to produce more output, then workers will have to be incentivized to work longer hours and wages and prices will increase. The bottom panel of the figure on the other hand represents the aggregate demand block of the economy: to achieve a given amount of real activity—output, again on the horizontal axis—real interest

[8]Central bank preferences are often described by a loss function that captures the idea that the central bank dislikes it when economic outcomes differ from the targeted outcomes. For example, the indifference curve we are plotting corresponds to the loss function

$$(y - y^*)^2 + (\pi - \pi^*)^2,$$

where y is output, π is inflation, and the starred variables are the targets for output and inflation. Unlike consumer theory where we maximize utility, here we want to minimize the loss function.

rates—nominal rates less inflation, shown on the vertical axis—need to be set at a certain level. The line is downward-sloping, reflecting the idea that higher interest rates depress aggregate demand, as discussed in Section 2.1.

The policy problem is to choose the nominal interest rate so as to minimize the deviations of output and inflation from their target values. As the policymaker changes the nominal interest rate, the real interest rate changes too, moving the economy along the aggregate demand curve (bottom panel) to determine the level of output. Moving to the upper panel, the aggregate supply curve determines the associated level of inflation. The solution to this optimal policy problem is straightforward: the monetary policymaker focuses on the upper panel of the figure and simply chooses the best feasible output-inflation pair. Visually, the optimal pair is given by the point of tangency with the indifference curve. The policymaker then uses the aggregate demand curve to determine which nominal interest rate to set in order to arrive at the desired level of output. The solution is shown by the dashed lines.

The role of inequality. How might household inequality affect this policy problem? Remember that we are assuming (for the moment) that inequality is not a target of the central bank, which implies that the policymaker indifference curve is not affected by inequality. Through the lens of the simple framework shown in Figure 5, most of the research on the connection between monetary policy and inequality discussed in the previous sections may be interpreted as studying ways in which household inequality could change the economy's aggregate demand relationship—that is, the mapping from interest rates to aggregate demand shown in the lower panel. However, the framing of the problem in the figure shows that changes in the demand block alone will not affect the optimal output and inflation outcomes, simply because the optimal policy choice is already fully pinned down by policymaker preferences and the supply side of the economy. It follows that changes in aggregate demand due to inequality will not affect the optimal inflation and output levels, though they may affect the nominal interest rate required to implement this optimal output-inflation allocation.

How big are those effects on optimal interest rates likely to be? Graphically, for a given output-inflation outcome, the optimal interest rate is determined by the *slope* and *intercept* of the economy's aggregate demand relationship. The slope of the curve reflects the sensitivity of aggregate demand to changes in real interest rates. As we discussed in Section 3, the heterogeneous-agent view of monetary policy transmission has not materially changed our broad understanding of this sensitivity. The intercept of the line, on the other hand, reflects forces that determine aggregate demand at any given interest rate, with changes in this

intercept reflecting so-called "aggregate demand shocks." Many plausible examples of such shocks are explicitly distributional in nature. For example, a tightening of credit conditions could require borrower households to reduce their debt levels, thus leading to a reduction in their spending and thus aggregate demand (e.g. Guerrieri & Lorenzoni, 2017). Alternatively, fiscal transfer payments to financially constrained households could lead to an increase in total consumer spending (e.g., Wolf, 2021). The simple theory sketched here suggests that such shifts in aggregate demand would lead to an equilibrium adjustment in interest rates while leaving optimal inflation and output outcomes unchanged.

4.2 A broad, distributional mandate

We now consider a central bank that explicitly incorporates distributional concerns as one of its policy goals, presumably along with its traditional output and inflation goals.[9] Long-run trends in inequality of course primarily reflect economic forces unrelated to monetary policy and the business cycle at large. However, short-run business-cycle fluctuations may well have material (short-term) effects on inequality, simply because aggregate shocks need not affect everyone in the same way. If a central bank's mandate includes distributional outcomes, then it will try to set its policy in a way that redistributes towards the hardest-hit households, essentially providing some insurance to those most exposed to aggregate shocks.

The role of insurance. In an ideal world, households would be able to buy insurance against all types of adverse events—including aggregate cyclical fluctuations—in private markets. For example, a worker could buy an insurance policy against the risk of unemployment. With such perfect insurance markets, standard macroeconomic models would predict that the consumption of all households would move up and down in proportion to aggregate consumption. Due to issues of moral hazard and adverse selection, however, many such insurance markets do not exist. With imperfect insurance markets, the economy will move away from this efficient pattern of risk sharing. Thus, following an aggregate shock, some households may be more severely affected than others and therefore reduce their consumption by more than the rest.

The social insurance benefit of filling in for these missing markets is widely recognized in other areas of public policy, including discussions of unemployment insurance systems,

[9]Studies of optimal monetary policy that incorporate distributional effects include Bhandari et al. (2021), Acharya et al. (2020), Le Grand et al. (2021), Dávila & Schaab (2022), and McKay & Wolf (2022). In the latter study we develop the views we describe in this section more formally.

tax policy, and social safety net programs. Similar underlying concerns may thus also guide optimal monetary policymaking: interest rates may be set in a way to both move aggregate consumption to the desired overall level *and* to smooth out consumption changes across households, essentially moving the cross-sectional consumption distribution closer to the desired efficient pattern of risk sharing.

Can monetary policy provide insurance? The evidence that we reviewed above is informative about the extent to which monetary policy can provide insurance and thus achieve such distributional objectives. Our main conclusion from that discussion was that monetary policy has rather evenly distributed effects across different groups of households—that is, expansionary monetary policy scales up the consumption of different households by similar proportions. To make the argument particularly stark, suppose for a moment that monetary policy was exactly distributionally neutral in the sense of scaling everyone's consumption up and down in perfect unison. In that case, monetary policy interventions would not bring us any closer to the efficient risk-sharing outcome, and so social insurance would not be a consideration for optimal monetary policy. Intuitively, even if a monetary policymaker would like to lean against inequality, monetary policy is not well-suited to do so, and so the monetary policymaker will act *as if* it had only a narrow mandate.

The actual situation is of course not as extreme as this—the effects of policy on consumption are not exactly equal in percentage terms—so there is some scope for monetary policy to alter the distribution of consumption. However, given the modest extent of these distributional effects, large changes in monetary policy would be needed to have a substantial effect on the consumption distribution. However, such large changes would likely be costly in terms of other policy goals (notably aggregate output and inflation stabilization). As a result, a central bank that targets both conventional aggregate outcomes as well as distributional outcomes is unlikely to deviate too much from the policies pursued by a central bank with a narrow mandate that just targets aggregate outcomes.

5 Conclusion

In this paper, we have taken stock of the recent research agenda that studies the connections between monetary policy and inequality, with three main conclusions. First, our reading of the empirical evidence suggests that monetary policy has a relatively uniform incidence across households. Second, accounting for micro heterogeneity across households changes our

understanding of the transmission channels for monetary policy. It has not, however, changed our understanding of the broad patterns of how monetary policy effects the macroeconomy. Third, our first two observations taken together somewhat limit the scope to which household inequality is likely to affect optimal monetary policy design, even if the central bank has a broad mandate that includes distributional considerations.

However, we emphasize that these broad conclusions come with several important qualifiers, which we view as topics for future research. First, in keeping with the recent academic literature on inequality and monetary policy, our analysis throughout this article focused on how heterogeneity interacts with the demand side of the economy. Changes to the supply side would affect optimal outcomes even for a central bank with a conventional narrow mandate. Second, our discussion has omitted the heterogeneity in consumption baskets across households. There is, however, evidence that low-income groups and racial minorities consume goods with more volatile prices (for example, Cravino et al., 2020; Lee et al., 2021), thus adding a further possible layer to the distributional effects of monetary policy. Finally, our conclusions on the distributional effects of monetary policy remain tentative, relying either on noisy consumption measures or assumptions on household consumption-savings decisions. More empirical work on these topics would be very welcome.

References

Acharya, S., Challe, E., & Dogra, K. (2020). Optimal monetary policy according to hank. CEPR Discussion Paper No. DP14429.

Amberg, N., Jansson, T., Klein, M., & Rogantini Picco, A. (2021). Five facts about the distributional income effects of monetary policy. CESifo Working Paper.

Andersen, A. L., Johannesen, N., Jørgensen, M., & Peydró, J.-L. (2021). Monetary policy and inequality. Working Paper.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, *109*(6), 2333–67.

Auclert, A., Rognlie, M., & Straub, L. (2018). The intertemporal keynesian cross. Technical report, National Bureau of Economic Research.

Bauer, M. D. & Swanson, E. T. (2022). A reassessment of monetary policy surprises and high-frequency identification. National Bureau of Economic Research Working Paper.

Beraja, M., Fuster, A., Hurst, E., & Vavra, J. (2019). Regional heterogeneity and the refinancing channel of monetary policy. *The Quarterly Journal of Economics*, *134*(1), 109–183.

Berger, D., Guerrieri, V., Lorenzoni, G., & Vavra, J. (2018). House prices and consumer spending. *The Review of Economic Studies*, *85*(3), 1502–1542.

Berger, D., Milbradt, K., Tourre, F., & Vavra, J. (2021). Mortgage prepayment and path-dependent effects of monetary policy. *American Economic Review*, *111*(9), 2829–78.

Berger, D. & Vavra, J. (2015). Consumption dynamics during recessions. *Econometrica*, *83*(1), 101–154.

Bernanke, B. S. & Kuttner, K. N. (2005). What explains the stock market's reaction to federal reserve policy? *The Journal of finance*, *60*(3), 1221–1257.

Best, M. C., Cloyne, J. S., Ilzetzki, E., & Kleven, H. J. (2020). Estimating the elasticity of intertemporal substitution using mortgage notches. *The Review of Economic Studies*, *87*(2), 656–690.

Bhandari, A., Evans, D., Golosov, M., & Sargent, T. J. (2021). Inequality, business cycles, and monetary-fiscal policy. *Econometrica*, *89*(6), 2559–2599.

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

Bloomberg. S&p 500 index quarterly closing values. (accessed Aug. 5, 2022).

Bureau of Economic Analysis (2022). National income and product accounts. (accessed Aug. 5, 2022).

Chang, M. & Schorfheide, F. (2022). On the effects of monetary policy shocks on earnings and consumption heterogeneity. CEPR Discussion Paper No. DP17049.

Chodorow-Reich, G., Nenov, P. T., & Simsek, A. (2021). Stock market wealth and the real economy: A local labor market approach. *American Economic Review*, *111*(5), 1613–57.

Cloyne, J., Ferreira, C., & Surico, P. (2020). Monetary policy when households have debt: new evidence on the transmission mechanism. *The Review of Economic Studies*, *87*(1), 102–129.

Coibion, O., Gorodnichenko, Y., Kueng, L., & Silvia, J. (2017). Innocent bystanders? monetary policy and inequality. *Journal of Monetary Economics*, *88*, 70–89.

Cravino, J., Lan, T., & Levchenko, A. A. (2020). Price stickiness along the income distribution and the effects of monetary policy. *Journal of Monetary Economics*, *110*, 19–32.

Dávila, E. & Schaab, A. (2022). Optimal monetary policy with heterogeneous agents: A timeless ramsey approach. Working Paper.

Di Maggio, M., Kermani, A., Keys, B. J., Piskorski, T., Ramcharan, R., Seru, A., & Yao, V. (2017). Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review*, *107*(11), 3550–88.

Doepke, M. & Schneider, M. (2006). Inflation and the redistribution of nominal wealth. *Journal of Political Economy*, *114*(6), 1069–1097.

Eichenbaum, M., Rebelo, S., & Wong, A. (2022). State-dependent effects of monetary policy: The refinancing channel. *American Economic Review*, *112*(3), 721–61.

Fagereng, A., Gomez, M., Gouin-Bonenfant, E., Holm, M., Moll, B., & Natvik, G. (2022). Asset-price redistribution. LSE Working paper.

Federal Reserve Board (2019). Survey of consumer finances. Accessed October 20, 2021.

Federal Reserve Board (2022). Finance rate on consumer installment loans at commercial banks. Consumer Credit - G.19 Release. (Accessed Aug 5, 2022).

Freddie Mac (2022). 30-year fixed rate mortgage average in the united states. (accessed Aug. 5, 2022).

Galí, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

Gertler, M. & Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, *7*(1), 44–76.

Greenwald, D. L., Leombroni, M., Lustig, H., & Van Nieuwerburgh, S. (2021). Financial and total wealth inequality with declining interest rates. Technical report, National Bureau of Economic Research.

Guerrieri, V. & Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics*, *132*(3), 1427–1467.

Guren, A. M., McKay, A., Nakamura, E., & Steinsson, J. (2021). Housing wealth effects: The long view. *The Review of Economic Studies*, *88*(2), 669–707.

Guvenen, F., Ozkan, S., & Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, *122*(3), 621–660.

Guvenen, F., Schulhofer-Wohl, S., Song, J., & Yogo, M. (2017). Worker betas: Five facts about systematic earnings risk. *American Economic Review*, *107*(5), 398–403.

Holm, M. B., Paul, P., & Tischbirek, A. (2021). The transmission of monetary policy under the microscope. *Journal of Political Economy*, *129*(10), 2861–2904.

Iacoviello, M. (2005). House prices, borrowing constraints, and monetary policy in the business cycle. *American economic review*, *95*(3), 739–764.

Institute for Social Research (2019). Panel study of income dynamics. Accessed July 28, 2022.

Johnson, D. S., Parker, J. A., & Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, *96*(5), 1589–1610.

Kaplan, G., Moll, B., & Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, *108*(3), 697–743.

Kaplan, G., Violante, G. L., & Weidner, J. (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*.

Le Grand, F., Martin-Baillon, A., & Ragot, X. (2021). Should monetary policy care about redistribution? optimal fiscal and monetary policy with heterogeneous agents. Technical report, Working Paper, SciencesPo.

Lee, M., Macaluso, C., & Schwartzman, F. (2021). Minority unemployment, inflation, and monetary policy. UCSD Working Paper.

Mavroeidis, S., Plagborg-Møller, M., & Stock, J. H. (2014). Empirical evidence on inflation expectations in the new keynesian phillips curve. *Journal of Economic Literature*, *52*(1), 124–88.

McKay, A. & Wieland, J. F. (2021). Lumpy durable consumption demand and the limited ammunition of monetary policy. *Econometrica*, *89*(6), 2717–2749.

McKay, A. & Wolf, C. K. (2022). Optimal policy rules in hank. MIT Working Paper.

Mian, A., Rao, K., & Sufi, A. (2013). Household Balance Sheets, Consumption, and the Economic Slump*. *The Quarterly Journal of Economics*, *128*(4), 1687–1726.

Mian, A., Straub, L., & Sufi, A. (2021). Indebted demand. *The Quarterly Journal of Economics*, *136*(4), 2243–2307.

Okun, A. M. (1973). Upward mobility in a high-pressure economy. *Brookings Papers on Economic Activity*, *4*(1), 207–262.

Patterson, C. (2022). The matching multiplier and the amplification of recessions. Working Paper.

Plagborg-Møller, M. & Wolf, C. K. (2021). Local projections and VARs estimate the same impulse responses. *Econometrica*, *89*(2), 955–980.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. *Handbook of macroeconomics*, *2*, 71–162.

Slacalek, J., Tristani, O., & 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.

Tenreyro, S. & Thwaites, G. (2016). Pushing on a string: Us monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, *8*(4), 43–74.

Tracy, J. & Wright, J. (2012). Why mortgage refinancing is not a zero-sum game.

Wolf, C. K. (2021). Interest rate cuts vs. stimulus payments: An equivalence result. National Bureau of Economic Research Working Paper 29193.

Wong, A. (2021). Refinancing and the transmission of monetary policy to consumption. Working paper.

Woodford, M. (2003). *Interest and prices: Foundations of a theory of monetary policy*. princeton university press.

A Details of indirect calculation of consumption effects

We report results from an indirect calculation of the consumption response to monetary policy. This calculation is premised on the following equation:

$$dc_i = \sum_j dp_j x_{ij} m_{ij} \tag{A.1}$$

where dc_i is the percentage change in consumption of household i after the change in monetary policy, dp_j is the change in the "price" of category j, x_{ij} is the exposure of household i to price j, and m_{ij} is the MPC of household i out of cash-flows of type j. We now describe how we construct each of these right-hand-side variables in turn.

Price changes. For each category j we construct an aggregate measure of how prices change. The data series are described in the table below. For housing returns we use the aggregate *value* of real estate held by US households, which in principle could reflect changes in quantities as well as prices. We are assuming that over a short period of time the change in value reflects price changes only.

We use the high-frequency monetary policy shocks identified by Gertler & Karadi (2015) to isolate quasi-random variation in monetary policy. For each variable, we run a simple recursive VAR with the shock ordered first (following the suggestion of Plagborg-Møller & Wolf, 2021), followed by real GDP, inflation, the 3-month Treasury bill rate, and the variable of interest. We use data from 1982Q3 to 2015Q3.[10] Using the estimated VAR we then construct impulse responses to the monetary policy shock. The average level of the impulse response over the first two years following the shock is our measure of the price effect of monetary policy. For mortgage rates and other interest rates, we subtract the inflation response to approximate a real interest rate.[11] Furthermore, we divide the price response of these two interest rates by 2 to account for the incomplete pass-through to the interest rates households actually pay (see Figure 2).[12]

[10] For business income, the results are very sensitive to including the Great Recession so for that variable we stop the sample at the end of 2007, but we note that this price change does not have a large impact on our overall calculation in any case.

[11] We measure the inflation response as the cumulative inflation response over the two years following the shock.

[12] The values reported in Table 1 are dp_j except for mortgages and other interest rates, where the values in the table are shown before the rescaling by 1/2.

Category	Data
Labor earnings	Wage and salary compensation, log real per capita $\log(COE) - \log(POP) - \log(GDPDEF)$ Bureau of Economic Analysis (2022)
Business income	Proprietors income plus dividend income, log real per capita $\log(PROPINC + B703RC1Q027SBEA) - \log(POP) - \log(GDPDEF)$ Bureau of Economic Analysis (2022)
Interest income	Interest income, log real per capita $\log(A064RC1Q027SBEA) - \log(POP) - \log(GDPDEF)$ Bureau of Economic Analysis (2022)
Return on stocks	S&P 500 index return, real $\log(SP500) - \log(GDPDEF)$ Bloomberg (2022)
Return on housing	Market value of real estate held by household sector, log real per capita $\log(HNOREMV) - \log(POP) - \log(GDPDEF)$ Bureau of Economic Analysis (2022) and Federal Reserve Board (2022)
Return on cash	Inflation measured by log change in GDP deflator $-4 \times \Delta \log(GDPDEF)$ Bureau of Economic Analysis (2022)
Mortgage rates	30-year fixed-rate mortgage, real $MORTGAGE30US/100$ less response of inflation as above Freddie Mac (2022)
Other interest rates	Rate on 4-year auto loan, real $TERMCBAUTO48NS/100$ less response of inflation as above Federal Reserve Board (2022)

Table A.1: Time series data description. All data apart from the S&P 500 were obtained from the Federal Reserve Bank of St Louis FRED database. The series identifiers are listed in the table.

HOUSEHOLD LEVEL EXPOSURES. We now describe how we construct the exposure of household i to price j. Our primary data source is the 2019 Survey of Consumer Finance (SCF). Take labor earnings for example. We set $x_{ij} = \gamma_{ij} y_{ij}/c_i$ where y_{ij} is wage and salary income of the household, c_i is an estimate of the baseline level of consumption, and γ_{ij} is the elasticity of the household's labor income with respect to changes in aggregate labor income. As the SCF does not include consumption data, we impute c_i using the household's income and the relationship between income and consumption expenditures in the 2019 wave of the Panel Study of Income Dynamics.[13] To understand this expression for x_{ij}, begin by supposing $\gamma_{ij} = 1 \forall i$. In this case, if aggregate wages increase by 1% (i.e. $dp_j = 0.01$), then $dp_j w_i$ represents a 1% increase in the wage income of household i. Dividing by c_i expresses this change in labor income as a proportion of their baseline consumption. Next, we incorporate the fact that low-income households are particularly exposed to business cycle fluctuations; i.e., γ_{ij} is decreasing in income. Specifically, earnings at the 10th percentile of the earnings distribution are about three times more sensitive to aggregate income than are earnings at the 80th percentile and the relationship is fairly linear in worker's rank in the earnings distribution below the 80th percentile (Guvenen et al., 2017). There is some evidence that earnings are more cyclical above the 80th percentile than at the 80th percentile, but this pattern did not hold in the Volcker disinflation (Guvenen et al., 2014), and therefore we assume a constant sensitivity above the 80th percentile. Overall we set $\gamma_{ij} = a - b \min(q_i, 0.8)$ where $q_i \in [0, 1]$ is the household's rank in the distribution of wage income. We choose b so that earnings are three times more sensitive at the 10th percentile than at the 80th percentile. The parameter a is normalized so that aggregate earnings have a unit elasticity to aggregate earnings. For business income and interest income we set $x_{ij} = y_{ij}/c_i$—that is we assume an equal exposure to the aggregate ($\gamma_{ij} = 1$).

Turning to asset categories, first take housing as an example. We set $x_{ij} = h_i/c_i$ where now h_i is the value of housing held by household i. If house prices increase by 1%, then $dp_j = 0.01$ and $dp_j x_{ij}$ will tell us how the home value changes relative to the household's baseline consumption level. Stocks and cash are treated similarly. Finally, for mortgages and other liabilities we set $x_{ij} = b_i/c_i$ where b_i is the debt outstanding. If the interest rate paid on the debt increases by dp_j then the debt service payment increases by $dp_j x_{ij}$ as a

[13] Specifically, we run a piecewise-linear regression of expenditure on income with a kink point at an income level of \$175,000. We used this piecewise-linear specification after examining the summary tables of the 2019 Consumer Expenditure Survey, which show expenditure rising nearly linearly in income up to the income level of \$150,000 to \$200,000. We experimented with adding richer household characteristics such as age, wealth and food consumption. These had little additional explanatory power in the PSID data.

share of baseline consumption. We include auto loans, credit card balances, and installment loans in the "other" debt category.

Marginal propensities to consume. We allow MPCs to vary across households and across cash-flow categories as described in Table 1. The variation across households is driven only by a classification of households into two groups: financially constrained and unconstrained. This classification is based on a threshold level of liquid assets. Households with liquid assets less than two weeks' income are categorized as constrained. Here liquid assets include transaction accounts and other liquid financial assets, but exclude non-financial assets and retirement assets. Following Slacalek et al. (2020), we use MPCs of 0.05 and 0.5 for unconstrained and constrained households' response to changes in income and the value of transaction accounts. For housing and stocks we use MPCs of 0.03, consistent with estimates from Guren et al. (2021) and Chodorow-Reich et al. (2021). If constrained households own these assets, then they are illiquid, and thus we use those low MPCs for all households. Lastly, we apply large MPCs to debt service payments. Theoretically, a change in a mortgage payment is a very persistent change in disposable income so we would expect a high MPC regardless of whether or not a household is constrained. Di Maggio et al. (2017) analyze the way households respond when their mortgage interest rates change. They find a large response of consumption as measured by auto purchases, but they also find that households on average use 8% of the increase in disposable income to pay down mortgage debt. Furthermore they find that high-income households are more likely to pay down debt. Motivated by this evidence, we set an MPCs of 0.8 for unconstrained households and 1.0 for constrained households.

Putting it all together and a consistency check. We construct consumption changes for each household according to equation (A.1). To construct Figure 3, we omit households with net worth in excess of $2.5 million, and then create quintiles of households by net worth, income, and age, and report the average dc_i within each quintile.[14]

The response of aggregate consumption implied by the figure is 0.9%, which encouragingly is similar to the 0.7% value that we obtain by including aggregate consumption in the monetary policy VAR that we used to construct the price changes.

[14]We chose to drop high-net worth households from this analysis because we are not confident that our consumption imputation and MPC assumptions apply well to these households.