

Five Facts about the Distributional Income Effects of Monetary Policy Shocks*

Niklas Amberg[†] Thomas Jansson[‡] Mathias Klein[§]

Anna Rogantini Picco[¶]

September 17, 2021

Abstract

We document five facts about the distributional income effects of monetary policy shocks using Swedish administrative individual-level data. (i) The effects of monetary policy shocks are U-shaped over the income distribution—i.e., expansionary shocks increase the incomes of high- and low-income individuals relative to middle-income individuals. (ii) The large effects in the bottom are accounted for by the labor-income response and (iii) those in the top by the capital-income response. (iv) The heterogeneity in the labor-income response is due to the earnings heterogeneity channel, whereas (v) that in the capital-income response is due to the income composition channel.

Keywords: Monetary policy; income inequality; heterogeneous agents; administrative data.

JEL: C55, E32, E52.

*For valuable comments and suggestions, we thank Pete Klenow (the editor), four anonymous referees, as well as Klaus Adam, Martin Blomhoff Holm, Mikael Carlsson, Fabio Canova, Cristiano Cantore, Marco Del Negro, Michael Ehrmann, Veronica Guerrieri, Tor Jacobson, Stefan Laséen, Benjamin Moll, Christian Moser (discussant), Ricardo Reis, Per Olof Robling, Kasper Roszbach, Jiri Slacalek, Ulf Söderström, Oskar Tysklind, and Erik Öberg, as well as conference and seminar participants at Banque de France, the Bank of Finland and CEPR Joint Webinar on New Avenues for Monetary Policy, the CEBRA 2021 Annual Meeting, the Bank of Canada Monetary Policy Research Workshop, the 11th ifo Conference on Macroeconomics and Survey Data, the Nordic Junior Macro Seminar Series, Norges Bank's 2020 Central Bank Macro Modelling Workshop, Norges Bank, Sveriges Riksbank, and the Tehran Institute for Advanced Study. The opinions expressed in this article are the sole responsibility of the authors and should not be interpreted as reflecting the views of Sveriges Riksbank.

[†]Sveriges Riksbank, Research Division. E-mail: niklas.amberg@riksbank.se.

[‡]Sveriges Riksbank, Research Division. E-mail: thomas.jansson@riksbank.se.

[§]Sveriges Riksbank, Research Division. E-mail: mathias.klein@riksbank.se.

[¶]Sveriges Riksbank, Research Division. E-mail: anna.rogantini.picco@riksbank.se.

1 Introduction

The distributional effects of monetary policy has become an important issue in monetary economics in recent years. There are two reasons for this. First, a growing literature on the effects of monetary policy in heterogeneous-agents New Keynesian (HANK) models suggests that micro-level heterogeneities are important drivers of the aggregate effects of monetary policy (see, e.g., Gornemann, Kuester and Nakajima, 2016; McKay, Nakamura and Steinsson, 2016; and Kaplan, Moll and Violante, 2018). Indeed, Auclert (2019) argues that redistribution is not merely a side-effect, but a *channel* through which monetary policy affects the real economy. Second, the rising levels of income inequality in most developed economies in recent decades have made distributional issues a key concern for the general public as well as for economic policymakers. These include central bankers, who have debated if and how monetary policy affects the distribution of incomes, and whether distributional effects should be taken into account in monetary policymaking (see, e.g., Mersch, 2014; Bernanke, 2015; and Draghi, 2016).

Determining the distributional effects is difficult, however, because monetary policy affects individuals' incomes through a large number of channels, many of which are likely to have opposite implications for the distribution of incomes (see Coibion et al., 2017, for an overview). Hence, to properly understand the distributional effects of monetary policy, one needs to determine not only its overall effects on the distribution of incomes, but also the respective roles of the different channels in driving the aggregate effect. The purpose of this paper is to contribute to such an understanding by presenting a set of new empirical facts about the heterogeneous individual-level income effects of monetary policy shocks.

Our empirical analysis is conducted on the basis of a monetary policy shock series identified using a state-of-the-art high-frequency approach (Nakamura and Steinsson, 2018; Jarociński and Karadi, 2020), and an administrative panel dataset comprising detailed, uncensored income data for every legal resident in Sweden over the period 1999-2018. We restrict the main estimation sample to prime working-age individuals, since a key objective of the paper is to inform the calibration of HANK-type models, which typically feature agents who do not retire.¹

¹Thus, all results reported in the main text of the paper concern 26-65 year old individuals. Corresponding results for a sample also comprising retirement-age individuals are reported in the Online Appendix.

We document five main facts about the distributional income effects of monetary policy shocks:

- (i) The total income effects of monetary policy shocks are U-shaped with respect to the income distribution. That is, expansionary shocks increase the incomes of low- and high-income individuals relative to middle-income individuals.
- (ii) The response of labor incomes to monetary shocks is large in the bottom of the distribution—which accounts for the strong total-income response of low-income individuals—but is small and statistically insignificant in the middle and top.
- (iii) The capital-income response to monetary shocks is statistically significant across the entire income distribution. The effect is particularly large in the very top, however, which drives the strong total-income response of high-income individuals.
- (iv) The heterogeneity in the labor-income response over the income distribution is accounted for by the *earnings heterogeneity channel*—that is, to a higher sensitivity of labor incomes to monetary shocks in the bottom than elsewhere in the distribution.
- (v) The heterogeneity in the capital-income response is, on the contrary, due to the *income composition channel*—that is, to the fact that capital income (i) responds particularly strongly to monetary policy shocks, and (ii) constitutes a larger share of total income for high-income individuals than for low- and middle-income individuals. The sensitivity of capital incomes to monetary shocks is, on the other hand, comparatively stable over the income distribution.

Two sets of remarks on the external validity of these facts are in order. The first is that they concern the effects of monetary policy *shocks*—i.e., surprise deviations from the systematic behavior of central banks—which typically account for a relatively small fraction of the actual changes in policy rates undertaken by central banks. Hence, our findings are directly informative about the distributional income effects of monetary policy in general insofar as the economy responds reasonably similarly to interest-rate changes induced by policy shocks as to changes driven by systematic policy. Importantly, though, studying the effects of monetary policy shocks is informative in its own

right, since doing so helps to constrain the class of models that can be expected to be good laboratories for studying systematic monetary policy.

The second set of remarks concern the relevance of the five facts for countries other than Sweden. Two important circumstances speak in favor of external validity in this regard. First, the Riksbank (the central bank of Sweden) conducts monetary policy on the basis of a modern inflation-targeting strategy and an institutional framework similar to those of, for example, the Federal Reserve, the ECB, and the Bank of England. Second, we provide evidence that the cross-sectional heterogeneity in the unconditional aggregate earnings risk of Swedish workers is similar to that of US workers documented by Guvenen et al. (2017), despite the many differences in labor-market institutions between the two countries. Thus, Sweden is a representative case in at least two dimensions relevant for the question at hand—namely, monetary policy and individuals' aggregate risk exposures. On the other hand, Sweden has a comparatively large social-welfare system and a high tax-to-GDP ratio, which likely reduces external validity somewhat, especially for the post-tax and post-transfer results.

Related literature. We contribute to the literature on the distributional income effects of monetary policy in several ways. In particular, our large-scale administrative data enables us to provide more detailed and precise results than most papers in the previous literature, which typically rely on survey data and/or time-series data on summary measures of income inequality, like the Gini coefficient and various percentile ratios (see, e.g., Coibion et al., 2017; Mumtaz and Theophilopoulou, 2017; Furceri, Loungani and Zdzienicka, 2018). The ability to study the very top of the income distribution—which is usually not possible with survey data, due to top-coding and measurement error—turns out to be important, since our results point to substantial heterogeneity in the effects of monetary policy shocks within the top of the distribution. Moreover, our results imply that the Gini coefficient is not well-suited for characterizing the distributional effects of monetary policy shocks, since the large effects observed in the tails of the income distribution mostly offset each other in this statistic. In general, our findings underscore the importance of considering the entire income distribution when studying the distributional effects of monetary policy.

Two contemporaneous papers also use administrative individual-level panel data to study distributional aspects of monetary policy shocks: Holm, Paul and Tischbirek

(2020) using Norwegian data and Andersen et al. (2020) using Danish data. The paper by Holm, Paul and Tischbirek (2020) differs in that it considers heterogeneous effects of monetary policy shocks along the liquid-asset distribution, whereas we focus on heterogeneity over the income distribution. We thus provide new and complementary empirical evidence relevant for HANK models, while also speaking more directly to the policy debate on the distributional income effects of monetary policy.

The paper by Andersen et al. (2020) is more similar in terms of question and empirical approach, but whereas we find that the income effects of expansionary monetary policy shocks are U-shaped over the income distribution, Andersen et al. find monotonically increasing effects. It is likely, however, that this difference at least partly can be explained by the fact that our main results concern prime working-age individuals, while Andersen et al. also include people above retirement age in their analysis. When we do the same, we obtain smaller income responses in the left tail of the income distribution—and thus a less pronounced U-shaped pattern—as retirees on average: (i) are less affected by monetary policy shocks; and (ii) have lower incomes than working-age individuals, and therefore disproportionately fall in the lower deciles of the income distribution.

The rest of this paper is structured as follows. Section 2 describes our data resources, specifies the econometric models, and explains the construction of the monetary policy shock series. Section 3 presents the results and Section 4 concludes.

2 Data and Empirical Framework

2.1 Econometric models

The empirical analysis consists of three main steps. First, we estimate how the effect of monetary policy shocks on individuals' total after-tax incomes varies over the income distribution. We do this using the following econometric model:

$$\frac{Y_{i,t+h}^T - Y_{i,t-1}^T}{Y_{i,t-1}^T} = \sum_{g=1}^{11} G_{i,t,g} \cdot \left[\alpha_g^{T,h} + \beta_g^{T,h} \cdot \widehat{\Delta i}_t \right] + \varepsilon_{i,t}^{T,h}, \quad (1)$$

which closely resembles the model used by Guvenen et al. (2017) to estimate individual-level unconditional earnings risk. The dependent variable is the percent

change in individual i 's real total after-tax income Y_i^T between years $t - 1$ and $t + h$; $G_{i,t,g} = \mathbf{1}_{G_{i,t}=g}$ is a binary indicator equal to one if individual i belongs to income group g in year t ; $\alpha_g^{T,h}$ is a group-specific intercept; and $\widehat{\Delta i_t}$ is the monetary policy shock in year t . $h = 0, 1, 2$ denotes the estimation horizon. Standard errors are two-way clustered at the individual and year levels, respectively, to account for within-individual serial correlation in the dependent variable (Bertrand, Duflo and Mullainathan, 2004) and within-year correlation in the monetary shock across individuals (Abadie et al., 2017). The coefficients of interest are the $\beta_g^{T,h}$, which capture the percent change in total after-tax income over an h -year horizon for individuals in income group g , following a contractionary monetary shock of one percentage point.

Second, to uncover the underlying drivers of the heterogeneities in the effects of monetary shocks on total after-tax incomes, we decompose the effects into the parts attributable to each component of total after-tax income. We conduct the decomposition exercise using the following model:

$$\frac{Y_{i,t+h}^k - Y_{i,t-1}^k}{Y_{i,t-1}^T} = \sum_{g=1}^{11} G_{i,t,g} \cdot \left[\alpha_g^{k,h} + \beta_g^{k,h} \cdot \widehat{\Delta i_t} \right] + \varepsilon_{i,t}^{k,h}, \quad (2)$$

which is identical to (1), except that the numerator in the dependent variable is the change in one of the components k in total after-tax income between years $t - 1$ and $t + h$, where $Y_{i,t}^T = \sum_k Y_{i,t}^k$. By constructing the dependent variables in this way and estimating (2) for each income component k , we obtain an exact decomposition of the estimated effects on total after-tax incomes into the effects attributable to each component—i.e., we then have $\beta_{g,h}^T = \sum_k \beta_{g,h}^k$. Thus, the contribution of component k to the effect of monetary policy shocks on total after-tax incomes is given by $\beta_{g,h}^k / \beta_{g,h}^T$.

Third, any heterogeneity in the effect of monetary policy on component k of total after-tax income is accounted for by some combination of (i) heterogeneity in the share of component k in total after-tax income (the income composition channel), and (ii) heterogeneity in the sensitivity of component k to monetary policy shocks (e.g., the earnings heterogeneity channel). To see this, note that the dependent variable in (2) can be rewritten as:

$$\frac{Y_{i,t+h}^k - Y_{i,t-1}^k}{Y_{i,t-1}^T} = \frac{Y_{i,t+h}^k - Y_{i,t-1}^k}{Y_{i,t-1}^k} \cdot \frac{Y_{i,t-1}^k}{Y_{i,t-1}^T}, \quad (3)$$

where the first factor is the percent change in income component k between years $t - 1$ and $t + h$, and the second is the share of component k in total after-tax income in year $t - 1$.

To disentangle the respective roles of these two sources of heterogeneity, we estimate (2) using a set of counterfactual dependent variables, in which the actual share of income component k in individual i 's total after-tax income—the second factor in (3)—is replaced by the corresponding average share in the entire sample. More specifically, we construct the counterfactual dependent variables by multiplying the original dependent variables by ξ_g , defined as:

$$\xi_g = \frac{Y_{g,t-1}^T}{Y_{g,t-1}^k} \cdot \frac{Y^k}{Y^T}, \quad (4)$$

where the first factor is the inverse of the average share of component k in total after-tax income in group g , and the second factor is the corresponding average share in the entire sample. The counterfactual dependent variables are thus defined as:

$$\frac{Y_{i,t+h}^k - Y_{i,t-1}^k}{Y_{i,t-1}^k} \cdot \underbrace{\frac{Y_{i,t-1}^k}{Y_{i,t-1}^T} \cdot \frac{Y_{g,t-1}^T}{Y_{g,t-1}^k} \cdot \frac{Y^k}{Y^T}}_{\approx 1} \approx \frac{Y_{i,t+h}^k - Y_{i,t-1}^k}{Y_{i,t-1}^k} \cdot \frac{Y^k}{Y^T}. \quad (5)$$

Hence, the counterfactual dependent variables can approximately be expressed as the product of the percent change in income component k for individual i between years $t - 1$ and $t + h$ and the average share of component k in total after-tax income across all individuals and years in the sample.² These variables thus capture what the growth in income component k —expressed as a share of total after-tax income—would have been if individual i 's income composition had been the same as the sample average. This exercise enables us to shut down the income composition channel and obtain a set of responses to monetary policy shocks in which any heterogeneity is due to differences in the sensitivity of a given component across the income distribution.

²The adjustment factor ξ_g only approximately cancels out the second term in (3) as it is defined at the group level, rather than at the individual level. The reason for not using the inverse of the *individual*-level income share in (4) is that we would then not be able to fully capture extensive-margin effects, such as when individuals go from zero labor income in year $t - 1$ to positive labor income in year $t + h$.

2.2 Data, sample, and variable definitions

The analysis is based on administrative register data from LISA, an annual panel comprising all legal residents in Sweden who are at least 16 year old. LISA is compiled by Statistics Sweden based on data from several official registries and is thus, unlike self-reported survey data, complete and largely free of measurement error. The data used in this paper is an extract from LISA which covers the period 1990–2018 and includes demographic variables, such as age and gender, as well as an individual’s total income and all its components. A key feature of the data is that the income variables are not top coded, which enables us to study income dynamics in the very top of the income distribution.

The main income concept in the analysis is total after-tax income, defined as the sum of labor income, capital income, and transfers, net of taxes. Labor income comprises earnings across all of an individual’s employers during a given year—including wages, salaries, bonuses, stocks and exercised stock options, bonds, and taxable employee benefits—as well as self-employment income. Capital income is the sum of net realized capital gains, dividends and interest income, and other capital income, net of interest expenses. Transfer income, finally, consists of a large number of components, including pension income, unemployment insurance, student grants, parental benefits, sickness and disability insurance, and incomes from job-training programs. The sum of labor income, capital income and transfers constitutes total pre-tax income. Total after-tax income is then obtained by subtracting income taxes owed from total pre-tax income.³ All income variables used in the analysis are deflated to real terms using the GDP price deflator with 2015 as base year.

We sort individuals into income groups g at an annual basis, using past average total *pre-tax* income as a proxy for permanent income. More specifically, in each year t , individual i ’s permanent income is computed as her average total pre-tax income in years $t - 4$, $t - 3$, and $t - 2$. When three years of past incomes are not available, we compute the average based on one or two years of data instead—an individual thus needs to be observed in years $t - 1$, $t + h$ and at least one year between $t - 4$ and $t - 2$ to be included in the sample. We then sort individuals into eleven income groups,

³We impute taxes owed for each individual-year observation based on the structure of the Swedish tax system in 2018. The details of the tax imputation are provided in Online Appendix B.

which correspond to deciles of the distribution of past average pre-tax income, except when it comes to the top decile, which we split into two: 90th to 99th and above the 99th, respectively. Note that there is no overlap between the periods over which income growth and past average income, respectively, are computed; hence, an individual's current income growth does not affect her current position in the income distribution.

We restrict the main sample to prime working-age individuals between 26 and 65 years old with positive total after-tax income, and the sample period to 1999-2018.⁴ To limit the influence of outliers, we drop observations for which the growth in total after-tax income exceeds 500 percent, or in which the growth in one of the main subcomponents of after-tax income—expressed as a share of after-tax income, as in (2)—exceeds 5 in absolute value. The resulting final sample comprises 79.5 million individual-year observations and 6.7 million unique individuals. Descriptive statistics for the main income variables and demographic characteristics by income group are provided in Table C1 in Online Appendix C.

2.3 The monetary policy shock series


We construct our monetary policy shock series, $\widehat{\Delta}_{it}$, by instrumenting changes in the repo rate—the Riksbank's main policy rate—with a monetary policy surprise series obtained from a high-frequency identification strategy similar to those used in the recent literature on monetary non-neutrality (see, e.g., Gertler and Karadi, 2015; Hanson and Stein, 2015; and Nakamura and Steinsson, 2018).

We define monetary surprises as changes in the yield of one-month Swedish Treasury bills on days of announcements of monetary policy decisions, adjusted for central bank information effects by means of Jarociński and Karadi's (2020) poor-man's sign restriction.⁵ This restriction involves setting the monetary surprise to zero in cases where

⁴We select 1999 as the start of our sample period because this was when the Riksbank's monetary policy decisions began to be communicated at regular and preannounced times, which is required for our high-frequency identification strategy to work. During our sample period, the Swedish economy experienced four recessions according to the standard OECD (2021) recession indicator (other business cycle dating methods tend to identify fewer recessions; cf. Edvinsson and Hegelund, 2018).

⁵An alternative would have been to define monetary surprises based on STINA contracts—overnight index swaps denominated in SEK—but our data on these contracts only begin in 2003. Reassuringly, our shock series is closely correlated with an analogously constructed series based on three-hour changes in STINA contracts around monetary announcements for the period 2003–2018, as shown in Figure C1 in Online Appendix C. The choice of daily changes in one-month T-bill rates as the basis of the monetary surprise series follows Flodén et al. (2021).

Figure 1: Construction of the monetary policy shock series



The data on the daily returns on the OMX Stockholm All Share Index, the yield on one-month Swedish Treasury bills, and the repo rate were all obtained from Sveriges Riksbank (2020). The data on monetary policy announcement dates were collected from Sveriges Riksbank (1999–2018).

stock returns on announcement days move in the same direction as the surprise in the market interest rate. More specifically, our surprise series comprises only those monetary policy announcements that fall in the second and fourth quadrants in Panel A of Figure 1, in which changes in the yield of one-month Swedish Treasury bills are plotted against the daily returns of the OMX Stockholm All Share Index on the days of monetary policy announcements.⁶ Observations in the first and third quadrants, on the

⁶On one announcement date (February 11, 2016), the one-month T-Bill rate exhibits a very large one-day swing, from -0.50 the day before, to -1.07 on the day of the announcement, and then back to -0.53 the day after. For this announcement date, we use the two-day change in the T-bill rate.

other hand, imply a positive comovement between interest rate changes and stock returns, which suggests that the surprise element in these announcements is due to news about the economy, rather than about the monetary stance. This motivates their exclusion from the monetary surprise series.

We then regress the change in the repo rate decided during monetary policy meeting Δi_m on the monetary surprise from the same meeting $\Delta i_m^{TBill'}$:

$$\Delta i_m = \alpha + \beta \cdot \Delta i_m^{TBill'} + \varepsilon_{i,t}. \quad (6)$$

Our basic monetary policy shock series, $\widehat{\Delta i_m}$, consists of the fitted values from the estimation of this regression (see Panel B of Figure 1 for a scatter plot illustrating the estimation). Finally, we aggregate the meeting-level shocks into an annual series by summing up all fitted values in a given year: $\widehat{\Delta i_t} = \sum_{m \in t} \widehat{\Delta i_m}$. Panels C and D of Figure 1 display the resulting shock series at monthly and annual frequency, respectively.

3 Results

This section presents the results of the empirical analysis. All reported results correspond to the effects of a 25 basis points expansionary monetary policy shock over a two-year horizon (results for estimation horizons $h = 0$ and $h = 1$ are provided in Online Appendix C). The reason for focusing on the two-year horizon here is that the peak of the effects that monetary shocks have on aggregate variables occurs about two years after the shock, as shown in Online Appendix A.⁷

3.1 The distributional income effects of monetary policy shocks

Panel A of Figure 2 reports the effects of an expansionary monetary policy shock on total after-tax incomes. The income groups are reported on the horizontal axis and the β^T , scaled by -0.25 , on the vertical axis.⁸ While monetary policy shocks have large and statistically significant effects on total after-tax incomes across the entire income

⁷More specifically, we show that an estimated proxy-VAR using our surprise series as an external instrument delivers aggregate impulse responses broadly in line with the textbook monetary policy transmission mechanism, where the peak effect on real activity occurs around two years after the shock.

⁸Corresponding results based on a sample that also includes retirement-age individuals are provided in Figure C3 in Online Appendix C.

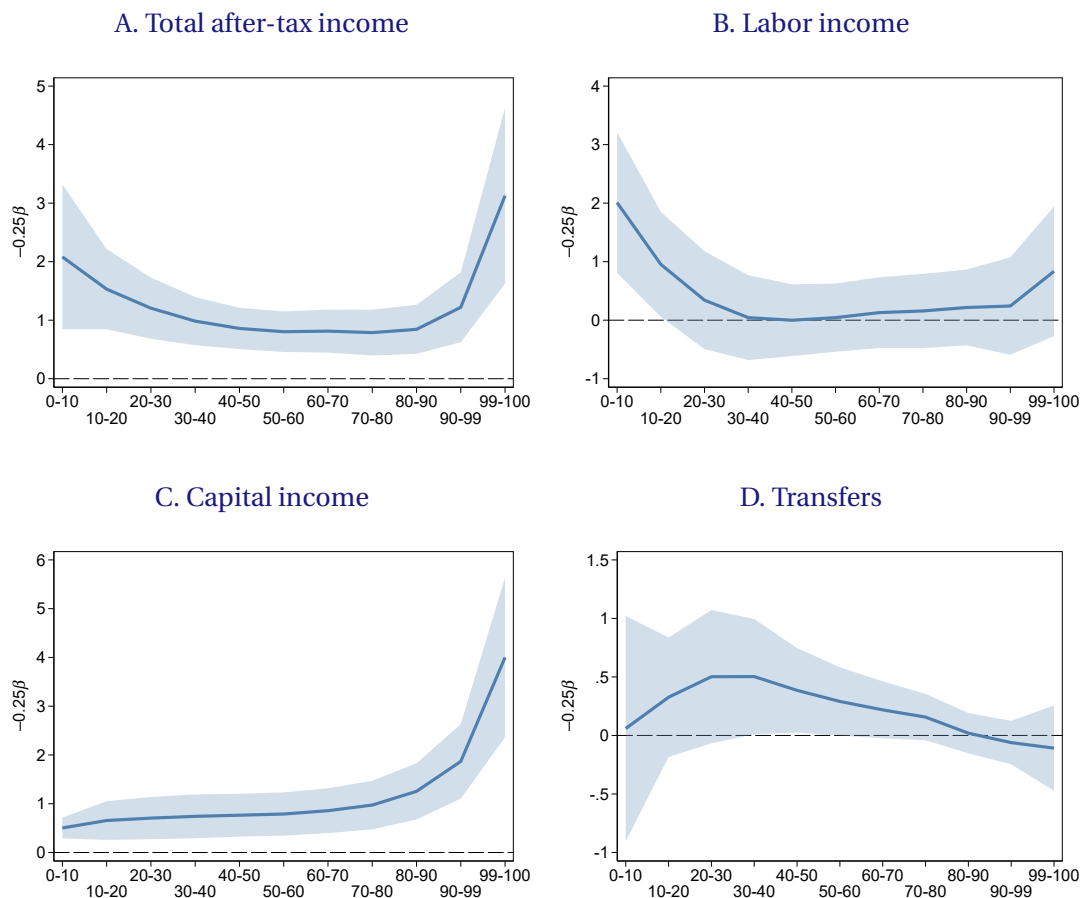
distribution, these effects are particularly large in the tails. More specifically, a 25 basis points reduction in the repo rate increases the after-tax incomes of the poorest and richest individuals by 2.1 and 3.1 percent, respectively, over a two-year horizon, whereas the corresponding response for middle-income individuals is 0.8 percent. Hence, the effects of monetary shocks on total after-tax incomes are around 3-4 times smaller in the middle of the distribution than in the tails, which yields a pronounced U-shaped pattern in the total-income response. Also, note that the total-income response is more than twice as large in the top percentile as in the rest of the top decile; hence, there is substantial heterogeneity within the top decile of the distribution.

Next, Panels B–D show the effects of monetary policy shocks on each of the three main components in total pre-tax income—labor income, capital income, and transfers. The response of labor income is large and statistically significant in the bottom two deciles, but statistically insignificant and mostly small throughout the rest of the distribution.⁹ The capital-income response, on the other hand, is statistically significant across the entire income distribution. The effect is particularly large in the very top—for example, the capital income response is around five times as large in the top percentile as in the middle of the distribution. The transfer response, finally, is hump-shaped with respect to the income distribution, but the estimated effects are mostly small and, with one exception, statistically insignificant in all income groups. The underlying drivers of the strong responses of total incomes in the top and bottom of the income distribution are thus different: labor income in the bottom and capital income in the top.

We end this subsection by addressing two potential concerns regarding the results reported in Figure 2. The first is that they could be unduly influenced by the monetary policy response to the Great Recession. Figure C2 in Online Appendix C shows, however, that the results are quite similar when the years 2007–10 are excluded from the sample. The second concern is that the results might be specific to our institutional setting. As a partial assessment of external validity, we undertake an exact replication of

⁹We do not observe hours and wages separately in the data and can therefore not provide direct evidence on the relative importance of these adjustment margins. However, when estimating equation (2) on a sample that only comprises continuously employed workers—defined as individuals whose labor income is at least half of the median annual labor income in Sweden in both $t - 1$ and $t + h$ —we obtain substantially smaller labor-income responses in the left tail of the distribution (see Figure C4 in Online Appendix C). Hence, the large labor-income effect in the left tail is primarily accounted for by individuals weakly attached to the labor market, and is thus most likely driven by changes in hours worked.

Figure 2: The effects of a -25bp shock on total after-tax income and its components



This figure shows the effects of a -25bp monetary policy shock on total after-tax income—as well as on the three main components of total pre-tax income—across the income distribution, as estimated using (1) and (2). The final component in after-tax income—taxes—is omitted from the figure for brevity (see row 6 of Table 2 for the estimated effects on taxes). The estimation horizon is $h = 2$. Shaded areas represent 90 percent confidence bands.

Guvenen et al.’s (2017) results on the systematic earnings risk of workers using Swedish data. The results, reported in Online Appendix D, shows that the cross-sectional patterns of (unconditional) earnings risk are very similar for Swedish and American workers, although the *levels* are generally higher for the latter. This speaks in favor of the relevance of our main results for other institutional settings, at least when it comes to labor income.

3.2 Inequality implications of the total-income effects

What do the total income effects of monetary shocks reported in Panel A of Figure 2 imply for aggregate income inequality? To answer this question, we undertake the following exercise. First, we compute the values of a number of commonly used measures of income inequality based on actual after-tax incomes in 2016 for all individuals in our sample. We then simulate the two-year effects of a –25 basis points monetary shock by multiplying each individual’s income in 2016 by $(1 - 0.25 \cdot \beta_{g,2}^T)$, where g is given by the income group to which an individual belongs in 2016. Finally, we compute the inequality measures for the simulated income distribution and compare the resulting values of the inequality measures with the initial values computed on actual data.

The results are reported in Table 1. The Gini coefficient changes very little after monetary policy shocks, as the large effects in the top and bottom mostly offset each other. We observe marked increases in the top income shares, however—especially in the top-1% share, which increases by almost two percent following a 25 basis points lowering of the repo rate. The increase in the ratio of the 90th to 50th percentile also points to a rise in income inequality following expansionary shocks, although the magnitude is very small. On the other hand, the standard deviation of log income as well as percentile ratios 90-10 and 50-10 all decrease, indicating that an expansionary monetary policy shock lowers income inequality.

These results imply that to fully understand the distributional consequences of monetary policy, one needs to look at its impact over the entire income distribution. The Gini coefficient, in particular, turns out to be problematic for characterizing the distributional effects of monetary policy, since it is virtually unaffected by monetary policy shocks, despite the pronounced heterogeneity in the individual-level effects.

3.3 Decomposing the total-income effects

How does each component in total after-tax income contribute to the total-income effects of monetary policy shocks? We presented some initial evidence in Panels B–D of Figure 2 that the large effects in the bottom of the distribution are driven by the labor-income response and the effects in the top by the response of capital incomes. To provide more detail and dig deeper into the drivers of the total-income effects, Table 2

Table 1: Implications of total income results for common measures of inequality

	Initial value	Two years after –25bp shock	Percent change
Gini coefficient	0.287	0.288	0.09
Top 1% income share	5.297	5.401	1.95
Top 10% income share	20.933	21.047	0.55
Standard deviation of log income	0.449	0.447	–0.43
Ratio of 90th to 10th percentile	3.226	3.212	–0.43
Ratio of 90th to 50th percentile	1.594	1.595	0.11
Ratio of 50th to 10th percentile	2.024	2.013	–0.55

This table reports the values of several common measures of income inequality computed based on actual total after-tax incomes in 2016 for all individuals in our sample (second column), as well as on a counterfactual income distribution, obtained by simulating the two-year effects of a –25 basis points monetary shock (third column). The rightmost column shows the percent change in the inequality measures after the simulated monetary policy shock.

reports the β^k , as before scaled by –25bp, from the estimation of (2) with each of the four main components in total after-tax income, as well as their respective subcomponents, as dependent variables.

The leftmost column shows that of the 2.1 percent increase in total after-tax incomes for the poorest individuals following a –25 basis points shock, 2.0 percentage points is due to labor income, 0.5 to capital income, 0.1 to transfer income, and –0.5 to taxes. In the middle of the distribution, capital incomes account for about two thirds of the total pre-tax income response and transfers for the remainder, as the contribution of labor income is close to zero; taxes then reduce the pre-tax income response by around 25 percent. For individuals in the top percentile, finally, the total after-tax income response is 3.1 percent, of which 4.0 percentage points is due to capital income, 0.8 to labor income, –0.1 to transfer income, and –1.6 to taxes.

Next, we decompose the labor-income response into the parts accounted for by wage income and self-employment income, respectively. Throughout the income distribution, the labor-income effects are driven entirely by the the wage-income response. The small contribution of self-employment income is explained by its very

Table 2: Decomposing the total income effects

	Income group										
	0-10	10-20	20-30	30-40	40-50	50-60	60-70	70-80	80-90	90-99	99-100
1. Labor income	2.0	1.0	0.3	0.0	0.0	0.0	0.1	0.2	0.2	0.2	0.8
1a. Wage income	1.9	0.9	0.3	0.1	0.0	0.1	0.2	0.2	0.2	0.2	0.8
1b. Self-empl. income	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2. Capital income	0.5	0.7	0.7	0.7	0.8	0.8	0.9	1.0	1.3	1.9	4.0
2a. Realized capital gains	0.3	0.3	0.4	0.4	0.4	0.5	0.5	0.6	0.8	1.2	3.0
2b. Dividends and interest	-0.1	-0.1	-0.1	0.0	0.0	0.0	0.0	0.0	0.1	0.2	0.6
2c. Interest expenses (-)	0.3	0.4	0.4	0.4	0.4	0.4	0.3	0.3	0.4	0.4	0.4
2d. Other capital income	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3. Market income	2.6	1.6	1.0	0.8	0.8	0.8	1.0	1.1	1.5	2.1	4.9
4. Transfer income	0.1	0.3	0.5	0.5	0.4	0.3	0.2	0.2	0.0	-0.1	-0.1
4a. Pensions	0.0	0.1	0.1	0.0	0.1	0.1	0.1	0.1	0.0	-0.1	-0.1
4b. Unemployment income	-0.4	-0.3	-0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4c. Other transfers	0.4	0.5	0.5	0.4	0.3	0.2	0.2	0.1	0.1	0.0	0.0
5. Total pre-tax income	2.6	2.0	1.5	1.3	1.1	1.1	1.2	1.3	1.5	2.1	4.8
6. Taxes (-)	-0.5	-0.4	-0.3	-0.3	-0.3	-0.3	-0.4	-0.5	-0.6	-0.8	-1.6
7. Total after-tax income	2.1	1.5	1.2	1.0	0.9	0.8	0.8	0.8	0.8	1.2	3.1

This table shows the contribution of each of the main components in total after-tax income—as well as of their respective subcomponents—to the total after-tax income effects of a -25bp monetary policy shock, as estimated using (1) and (2). The estimation horizon is $h = 2$. Two income components (interest expenses and taxes) enter after-tax income with a negative sign; for these components, a negative number in the table implies that the component *increases* after an expansionary monetary policy shock, and thus contributes to a *decrease* in after-tax income, and vice versa.

small average share in labor income over most of the income distribution (see Table C1 in Online Appendix C).¹⁰

The capital-income response is decomposed into the parts due to realized capital

¹⁰Note, however, that self-employment income in the official Swedish income statistics only comprises income from self-proprietorships and trading partnerships; the incomes of individuals who are self-employed in incorporated firms are instead classified as wage income or dividends. Hence, our data likely understates the role of self-employment income.

gains, dividends and interest, interest expenses, and other capital income. Realized capital gains and interest expenses jointly account for the entire capital-income effect in all but the top two income groups, where dividends and interest income contribute a small, but non-negligible part to the total. In the bottom half of the distribution, realized capital gains and interest expenses are equally important, but the relative importance of the former grows monotonically over the top half of the distribution—in the top percentile, the contribution of realized capital gains to the total capital-income effect is more than seven times as large as that of interest expenses.¹¹

Finally, as the transfer-income response is statistically insignificant across almost the entire income distribution, we refrain from drawing conclusions about the respective contributions of its subcomponents.

3.4 Income composition versus within-component heterogeneity

Heterogeneity in the effects of monetary policy shocks on component k of total after-tax income is, as discussed in Section 2.1, accounted for by some combination of heterogeneity in the share of component k in total after-tax income—the *income composition channel*—and heterogeneity in the sensitivity of component k to monetary policy shocks. We refer to the latter as *within-component heterogeneity*, but we follow the terminology of Coibion et al. (2017) when considering specific income components; for example, we refer to within-component heterogeneity in labor income as the *earnings heterogeneity channel*.

Before turning to the formal analysis of the respective roles of these two channels, it is useful to present some descriptive statistics on how the composition of incomes varies over the distribution. This is done in Panel A of Figure 3.¹² With the exception of the first income group, labor income constitutes the largest share of total pre-tax income and is inversely U-shaped over the income distribution—that is, labor income is relatively less important in the bottom and top than in the middle. The share of

¹¹In Online Appendix E, we analyze how the values of individual-level asset holdings respond to monetary policy shocks. The results indicate that the gains in housing values are fairly equally distributed across the income distribution, but that the gains in stock values are strongly concentrated in the right tail, and in particular in the top percentile.

¹²For ease of interpretation, we report the statistics on income composition in Panel A in terms of shares of *pre-tax* incomes, even though the counterfactual estimates in Panels B-D are constructed based on shares of *after-tax* incomes.

capital income is mildly negative over the first nine deciles of the income distribution, but then increases sharply, reaching 16 percent in the top percentile.¹³ The share of transfer income, finally, is large in the bottom (52 percent of total pre-tax income in the first decile), but then decreases monotonically over the income distribution.

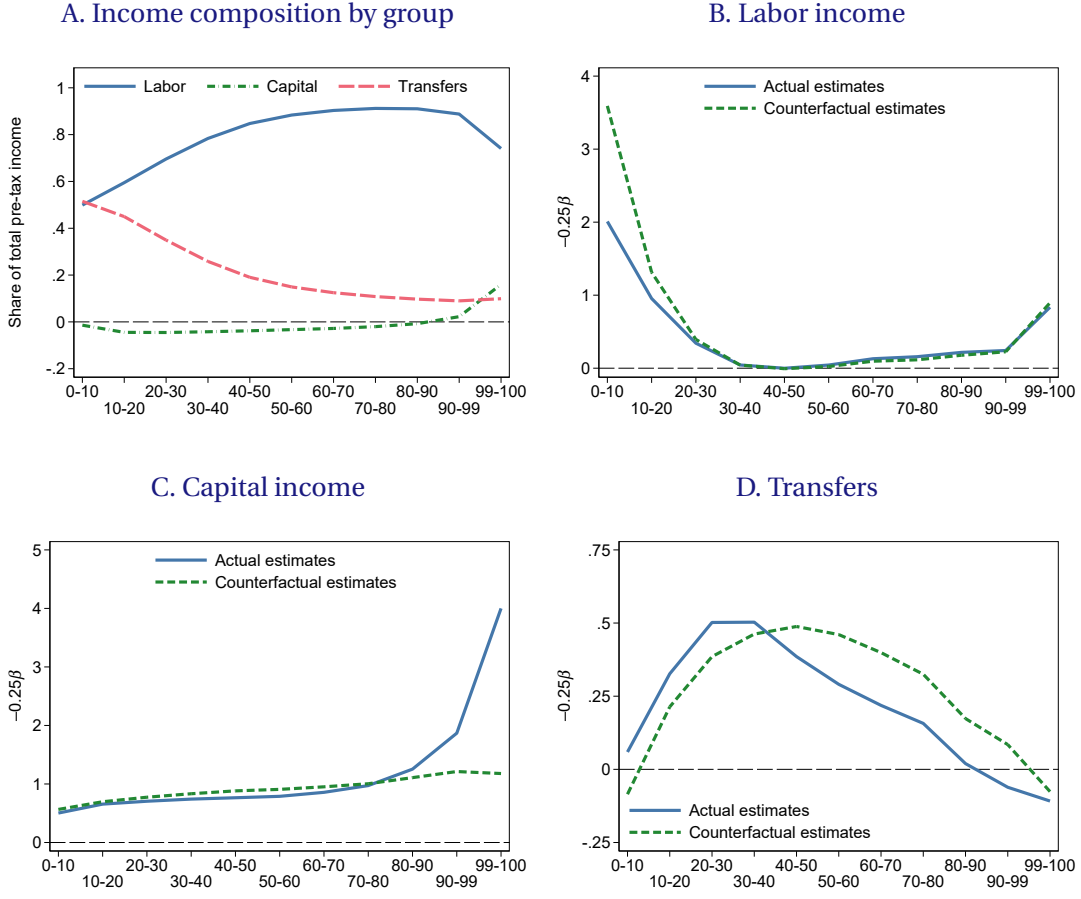
Turning to the formal analysis, we compare the actual estimates of β^k —where both the within-component heterogeneity and the income composition channels are operative—with the counterfactual estimates, where the income composition channel is shut down. This enables us to assess the respective roles of these two channels in accounting for the observed heterogeneity over the income distribution in the responses of each income component k to monetary policy shocks. The results are reported in Panels B–D of Figure 3. The solid blue lines are the actual estimates reported in Figure 2, while the dashed green lines are the counterfactual estimates described in Section 2.1.

Consider first the labor-income results, reported in Panel B. The counterfactual estimates are considerably larger than the actual estimates in the bottom of the distribution, but virtually equivalent in the middle and top. This implies that the heterogeneity in the labor-income effects of monetary policy shocks—namely, strong effects in the bottom of the distribution, but small and insignificant effects in the middle and top—is accounted for by the earnings heterogeneity channel. That is, the heterogeneity is due to the fact that labor incomes are more sensitive to monetary shocks in the bottom of the income distribution than in the middle and top.

The capital-income results are reported in Panel C. While the actual estimates show strong heterogeneity in the capital-income response—with the effects in the top being around five times larger than in the middle—the counterfactual estimates are comparatively stable over the income distribution. This implies that the heterogeneity in the capital-income response is due to the income composition channel—that is, to the fact that (i) capital income responds particularly strongly to monetary policy shocks, and (ii) capital income constitutes a larger share of total incomes in the top of the distribution. Hence, our results suggest that the various channels that may generate heterogeneity in the sensitivity of capital incomes to monetary shocks over the income

¹³In the deciles with negative capital-income shares, interest expenditures are on average larger in absolute value than all other capital income components taken together.

Figure 3: Income composition versus within-component heterogeneity



Panel A plots the average shares of labor income, capital income, and transfers, respectively, in total pre-tax income by income group. In Panels B-D, the solid blue lines correspond to the actual effects of a -0.25β monetary policy shock for estimation horizon $h = 2$, as estimated using (2), while the dashed green lines are the responses obtained when estimating (2) with the counterfactual dependent variables described in Section 2.1. More specifically, for each main income component (labor income, capital income, and transfers), the counterfactual estimate is obtained by first computing the counterfactual estimates for its respective subcomponents and then adding these up.

distribution—such as the savings redistribution, financial segmentation, and portfolio channels (Coibion et al., 2017)—are relatively unimportant in quantitative terms.

Panel D, finally, shows that the actual and counterfactual estimates for transfer incomes track each other fairly closely over the income distribution. We again refrain from drawing conclusions based on the transfer-income results, as the estimated effects are mostly statistically insignificant and of small magnitudes. In sum, the hetero-

geneity in the labor-income response is accounted for by the earnings heterogeneity channel, while the heterogeneity in the capital-income response is due to the income composition channel.

4 Conclusions

This paper has presented a set of new empirical facts about the distributional income effects of monetary policy shocks. In particular, we have shown that the effects of monetary policy shocks on individuals' total after-tax incomes are U-shaped with respect to the income distribution—that is, expansionary monetary shocks increase the incomes of low- and high-income individuals relative to middle-income individuals. The U-shaped response is, in turn, the result of a strong labor-income response in the bottom of the distribution and a strong capital-income response in the top.

Our analysis provides a basis for several interesting areas of future research. First, as our analysis has focused entirely on conventional (interest-rate based) monetary policy, a natural extension would be to consider the distributional consequences of unconventional monetary interventions, like the asset-purchase programs that many central banks have undertaken during the last decade. Second, our empirical framework, based on large-scale individual-level administrative data, may also be used to provide insights on the distributional effects of, for example, fiscal policy and macroprudential interventions.

References

- Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey Wooldridge.** 2017. “When Should You Adjust Standard Errors for Clustering?” NBER Working Paper No. 24003.
- Andersen, Asger Lau, Niels Johannesen, Mia Jørgensen, and José-Luis Peydró.** 2020. “Monetary Policy and Inequality.” Department of Economics and Business, Universitat Pompeu Fabra Economics Working Papers 1761.
- Auclert, Adrien.** 2019. “Monetary Policy and the Redistribution Channel.” *American Economic Review*, 109(6): 2333–67.
- Bernanke, Ben S.** 2015. “Monetary Policy and Inequality.” The Brookings Institution. Available at: <https://www.brookings.edu/blog/ben-bernanke/2015/06/01/monetary-policy-and-inequality/>.
- Bernanke, Ben S., and Kenneth N. Kuttner.** 2005. “What Explains the Stock Market’s Reaction to Federal Reserve Policy?” *Journal of Finance*, 60(3).
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan.** 2004. “How Much Should We Trust Differences-In-Differences Estimates?” *Quarterly Journal of Economics*, 119(1): 249–275.
- Calvet, Laurent E., John Y. Campbell, and Paolo Sodini.** 2007. “Down or Out: Assessing the Welfare Costs of Household Investment Mistakes.” *Journal of Political Economy*, 115(5): 707–747.
- Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia.** 2017. “Innocent Bystanders? Monetary Policy and Inequality.” *Journal of Monetary Economics*, 88(C): 70–89.
- Draghi, Mario.** 2016. “Stability, Equity, and Monetary Policy.” 2nd DIW Europe Lecture. Available at: <https://www.ecb.europa.eu/press/key/date/2016/html/sp161025.en.html>.
- Edvinsson, Rodney, and Erik Hegelund.** 2018. “The Business Cycle in Historical Perspective. Reconstructing Quarterly Data on Swedish GDP, 1913-2014.” *Journal of European Economic History*, 2018(1): 33–60.
- Flodén, Martin, Matilda Kilström, Jósef Sigurdsson, and Roine Vestman.** 2021. “Household Debt and Monetary Policy: Revealing the Cash-Flow Channel.” *Economic Journal*, 131(636): 1742–1771.

- Furceri, Davide, Prakash Loungani, and Aleksandra Zdzienicka.** 2018. “The Effects of Monetary Policy Shocks on Inequality.” *Journal of International Money and Finance*, 85(C): 168–186.
- Gertler, Mark, and Peter Karadi.** 2015. “Monetary Policy Surprises, Credit Costs, and Economic Activity.” *American Economic Journal: Macroeconomics*, 7(1): 44–76.
- Gornemann, Nils, Keith Kuester, and Makoto Nakajima.** 2016. “Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy.” Board of Governors of the Federal Reserve System International Finance Discussion Papers No. 1167.
- Güvenen, Fatih, Sam Schulhofer-Wohl, Jae Song, and Motohiro Yogo.** 2017. “Worker Betas: Five Facts about Systematic Earnings Risk.” *American Economic Review*, 107(5): 398–403.
- Hamermesh, Daniel S.** 2007. “Viewpoint: Replication in Economics.” *Canadian Journal of Economics*, 40(3): 715–733.
- Hanson, Samuel G., and Jeremy C. Stein.** 2015. “Monetary Policy and Long-Term Real Rates.” *Journal of Financial Economics*, 115(3): 429–448.
- Holm, Martin, Pascal Paul, and Andreas Tischbirek.** 2020. “The Transmission of Monetary Policy under the Microscope.” *Journal of Political Economy*, forthcoming.
- Jarociński, Marek, and Peter Karadi.** 2020. “Deconstructing Monetary Policy Surprises—The Role of Information Shocks.” *American Economic Journal: Macroeconomics*, 12(2): 1–43.
- Jentsch, Carsten, and Kurt G. Lunsford.** 2019. “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States: Comment.” *American Economic Review*, 109(7): 2655–78.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante.** 2018. “Monetary Policy According to HANK.” *American Economic Review*, 108(3): 697–743.
- Laséen, Stefan.** 2020. “Monetary Policy Surprises, Central Bank Information Shocks, and Economic Activity in a Small Open Economy.” Sveriges Riksbank Working Paper No. 396.
- Lundberg, Jacob.** 2019. “Framtidens skatter: En reform för tillväxt och enkelhet.” Technical report. Available at: <https://timbro.se/app/uploads/2019/11/framtidens-skatter.pdf>.
- McKay, Alisdair, Emi Nakamura, and Jón Steinsson.** 2016. “The Power of Forward Guidance Revisited.” *American Economic Review*, 106(10): 3133–58.

- Mersch, Yves.** 2014. “Monetary Policy and Economic Inequality.” Keynote speech. Available at: https://www.ecb.europa.eu/press/key/date/2014/html/sp141017_1.en.html.
- Mumtaz, Haroon, and Angeliki Theophilopoulou.** 2017. “The Impact of Monetary Policy on Inequality in the UK. An Empirical Analysis.” *European Economic Review*, 98(C): 410–423.
- Nakamura, Emi, and Jón Steinsson.** 2018. “High-Frequency Identification of Monetary Non-Neutrality: The Information Effect.” *Quarterly Journal of Economics*, 133(3): 1283–1330.
- OECD.** 2021. “Composite leading indicators.” OECD, Main Economic Indicators (accessed August 25, 2021).
- Rigobon, Roberto, and Brian Sack.** 2004. “The Impact of Monetary Policy on Asset Prices.” *Journal of Monetary Economics*, 51(8): 1553–1575.
- Statistics Sweden.** 2021. “Månadslön, Median: Samtliga Utbildningsnivåer, totalt (AM0112C5).” Statistics Sweden, Statistikdatabasen (accessed September 4, 2021).
- Sveriges Riksbank.** 1999–2018. “Monetary Policy Minutes.” Various issues. Available at: <https://www.riksbank.se/en-gb/press-and-published/minutes-of-the-executive-boards-monetary-policy-meetings/>.
- Sveriges Riksbank.** 2020. “DORIS – The Riksbank Data Library.”
- Swedish Ministry of Finance.** 2018. “Beräkningskonventioner.” Technical report.
- Waldenström, Daniel.** 2014. “Swedish Stock and Bond Returns, 1856–2012.” In *Historical Monetary and Financial Statistics for Sweden, vol. 2*. Ed. Rodney Edvinsson, Tor Jacobson and Daniel Waldenström. Stockholm:Sveriges Riksbank and Ekerlids förlag.
- Williams, John C.** 2015. “Measuring Monetary Policy’s Effect on House Prices.” Federal Reserve Bank of San Francisco Economic Letter No. 2015-28.

**Online Appendix for “Five Facts About the Distributional
Income Effects of Monetary Policy Shocks”**


Niklas Amberg, Thomas Jansson, Mathias Klein, and Anna Rogantini Picco

Appendix A. Aggregate Effects of the Monetary Policy Shock

In this section, we validate our monetary policy shock series by estimating a proxy-VAR and studying the induced aggregate dynamics. More specifically, we use the monthly monetary policy surprise series described in Section 2.3 as an external instrument in a VAR that includes the following variables: the repo rate, the log of industrial production, the unemployment rate, and a measure of underlying inflation published by Sveriges Riksbank. The VAR includes 12 lags, a constant, and a linear time trend, and the estimation strategy follows Gertler and Karadi (2015) and Jarociński and Karadi (2020). Moreover, we use the moving blocks bootstrap that has recently been recommended by Jentsch and Lunsford (2019) for proxy-VARs in order to appropriately take into account the uncertainty about the relation between the structural shocks and the instruments and thus to obtain consistent confidence bands. The first-stage F-statistic is 9.02, which implies that weak-instrument problems are unlikely to be a major concern for the analysis.

Figure A1 shows the results of the proxy-VAR where we normalize the impulse responses such that the repo rate falls by 25 basis points in the impact period. The lightly and darkly shaded areas indicate 90 and 68 percent confidence bands, respectively, obtained from 1,000 bootstrap repetitions. The exogenous fall in the repo rate leads to a significant increase in real economic activity with a peak response after around two years. After a mild increase in the first periods, the unemployment rate falls and then slowly converges back to its pre-shock level. In addition, inflation increases already on impact and shows a positive response until the end of the forecast horizon. Overall, these responses are broadly in line with the standard monetary policy transmission mechanism at the aggregate level, which validates the use of our monetary policy shock series when studying the individual-level effects of monetary policy shocks.

Figure A1: Proxy VAR



Solid lines show point estimates in response to an exogenous fall in the repo rate by 25 basis points in the impact period. The darkly and lightly shaded areas correspond to 68 and 90 percent bootstrapped confidence intervals, respectively. The unit of the horizontal axis is a month and the sample is 1999M1-2018M12.

Appendix B. Details on the Imputation of Taxes Owed

We impute taxes owed for each individual-year observation in our sample on the basis of the structure of the Swedish tax system in 2018. As labor income and capital income are taxed separately, we first explain how each component is computed, and then how they are added together.

Labor income taxation in Sweden takes its starting point in assessed earned income (*fastställd förvärvsinkomst*), which is obtained by summing labor income, self-employment income, unemployment benefits, and pension income. A basic deduction (*grundavdrag*) is then subtracted from the assessed earned income, which yields taxable earned income (*beskattningsbar inkomst*). The labor income taxes owed is then computed on the basis of taxable earned income as follows. First, taxable earned income in its entirety is taxed according to the municipal tax rate in the individual's municipality of residence. Second, any part of taxable earned income exceeding a certain threshold (*brytpunkt 1*) is subject to an additional 20 percentage points in central-government tax. Third, any part of taxable earned income exceeding a second, higher threshold (*brytpunkt 2*) is subject to an additional five percentage points in central-government tax. Finally, an earned-income tax credit (*jobbskatteavdrag*) is deducted from the sum of municipal and central-government taxes owed to obtain the final labor income taxes owed. The respective sizes of the basic deduction and the earned-income tax credit are determined by formulas in which the arguments are the assessed earned income, the basic deduction, and the PBB (basic price amount, or *prisbasbelopp*).^{B1}

When computing labor income taxes owed, we use the average municipal tax rate in 2018 (32.93 percent) throughout the entire sample period. For the computation of the basic deduction and the earned-income tax credit, we use the actual PBB in every year, but follow the formulas determining their respective sizes as of 2018. Similarly, we use the actual thresholds for central-government tax in every year—these thresholds are automatically increased every year by the inflation rate plus two percentage points, unless parliament actively decides on some other change. In this way, we compute labor income taxes owed in a way that is consistent over time, while avoiding drift in

^{B1}The PBB is a reference number used in various parts of the Swedish public-finance system and changes every year in line with the consumer price index. The formulas used for computing the basic deduction and the earned-income tax credit in 2018 are specified in Tables 3.4 and 3.10 in Swedish Ministry of Finance (2018).


the shares of individuals falling in the various tax brackets.

Capital income is generally taxed at a flat rate of 30 percent in Sweden. The effective tax rate is lower for a few important asset classes, however, as only a fraction of the income from these classes is subject to taxation—the effective tax rates on capital incomes therefore vary between 20 and 30 percent depending on asset class. We do not observe the incomes from all asset classes separately and can therefore not apply the exact effective tax rate for the incomes from each asset class when imputing capital income taxes. Instead, we use the weighted average of the effective tax rates across all asset classes—which amounts to 23.5 percent—in the computations (Lundberg, 2019, Table 12). Interest expenses give rise to deductions, which amount to 30 percent for expenses up to 100,000 SEK, and to 21 percent for expenses exceeding 100,000 SEK.

Labor and capital taxes are then summed to obtain total income taxes. In cases where capital income taxes are negative—i.e., when capital losses and interest expenses exceed capital gains and other capital incomes—the negative amount is deducted from labor income taxes owed. Note, however, that the sum of labor and capital taxes is not allowed to be negative; hence, when capital taxes are negative and exceed labor taxes in absolute value, total income taxes owed are set to zero.

Figure B1 plots the average tax rate—defined as total income taxes divided by total pre-tax income—by income group for the sample period 1999–2018. The progressivity of the tax system is evident, with the average tax rate increasing monotonically over the income distribution, from around 10 percent in the bottom decile to around 40 percent in the top percentile.

Figure B1: Average tax rate by income group



This figure shows the average tax rate—defined as total income taxes divided by total pre-tax income—by income group for the sample period 1999–2018. See the main text in Online Appendix B for details on the imputation of income taxes.


Appendix C. Additional Tables and Figures

This appendix provides additional tables and figures referred to in the main text of the paper. Table C1 presents descriptive statistics by income group; Figure C1 a comparison of our monetary shock series with an analogously constructed series based on STINA contracts; Figure C2 the total after-tax income responses when the financial crisis is excluded from the estimation sample; Figure C3 the total after-tax income responses when retirement-age individuals are included in the sample; Figure C4 the labor-income responses when only continuously employed individuals are included in the sample; and Figure C5 the effects of a -25bp monetary policy shock on total after-tax income and the three main components of pre-tax incomes for estimation horizons $h = 0$ and $h = 1$.

Table C1: Descriptive statistics by income group


	Income group										
	0-10	10-20	20-30	30-40	40-50	50-60	60-70	70-80	80-90	90-99	99-100
A. Total income											
Average pre-tax income (1,000s)	144	196	226	252	278	305	336	376	445	651	1,852
Average after-tax income (1,000s)	123	164	185	204	221	239	259	285	323	427	1,144
B. Average shares of total pre-tax income											
Labor income	0.50	0.59	0.70	0.78	0.85	0.88	0.90	0.91	0.91	0.89	0.74
- <i>Wage income</i>	0.41	0.54	0.66	0.75	0.82	0.86	0.88	0.89	0.89	0.87	0.72
- <i>Self-employment income</i>	0.09	0.05	0.04	0.03	0.02	0.02	0.02	0.02	0.02	0.02	0.02
Capital income	-0.01	-0.04	-0.05	-0.04	-0.04	-0.03	-0.03	-0.02	-0.01	0.02	0.16
- <i>Realized capital gains</i>	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.03	0.07
- <i>Dividends and interest income</i>	0.03	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.04	0.14
- <i>Interest expenses</i>	-0.07	-0.07	-0.07	-0.07	-0.06	-0.06	-0.06	-0.05	-0.05	-0.05	-0.05
- <i>Other capital income</i>	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.01
Transfer income	0.52	0.45	0.35	0.26	0.19	0.15	0.12	0.11	0.10	0.09	0.10
- <i>Pension income</i>	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.04	0.05	0.08
- <i>Unemployment benefits</i>	0.07	0.08	0.08	0.06	0.04	0.03	0.02	0.01	0.01	0.01	0.01
- <i>Other transfer income</i>	0.42	0.34	0.24	0.17	0.12	0.09	0.07	0.06	0.05	0.03	0.02
Taxes	-0.11	-0.14	-0.16	-0.18	-0.20	-0.21	-0.22	-0.24	-0.26	-0.33	-0.39
C. Other characteristics (means)											
Male	0.41	0.34	0.32	0.36	0.42	0.51	0.59	0.64	0.68	0.75	0.84
Age	41.9	42.5	43.5	44.4	45.2	45.8	46.3	47.2	48.1	49.8	52.0
Less than high-school education	0.28	0.20	0.17	0.17	0.16	0.16	0.15	0.14	0.11	0.06	0.05
High-school education	0.41	0.48	0.54	0.56	0.57	0.54	0.51	0.47	0.40	0.29	0.22
Post-secondary education	0.31	0.32	0.28	0.27	0.27	0.30	0.33	0.39	0.49	0.65	0.73

Figure C1: Comparison of monetary policy shock series (1M T-bill versus STINA)




This figure compares the monetary policy shock series used in the paper (solid blue line) with an analogously constructed series based on three-hour changes in overnight index swap rates around monetary announcements (STINA contracts; dashed green line) over the period for which we have data on STINA contracts. The data on STINA surprises are from Laséen (2020).

Figure C2: Total after-tax income results with and without financial crisis ($h = 2$)



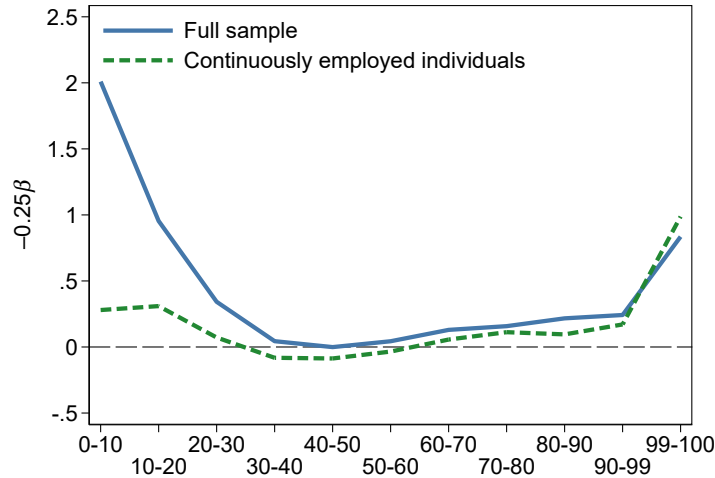
This figure shows the effects of a -25bp monetary policy shock on total after-tax incomes, as estimated using (1), when the financial-crisis years (2007–10) are included and excluded, respectively, from the estimation sample. The estimation horizon is $h = 2$.

Figure C3: Total after-tax income results without and with people above 65 ($h = 2$)



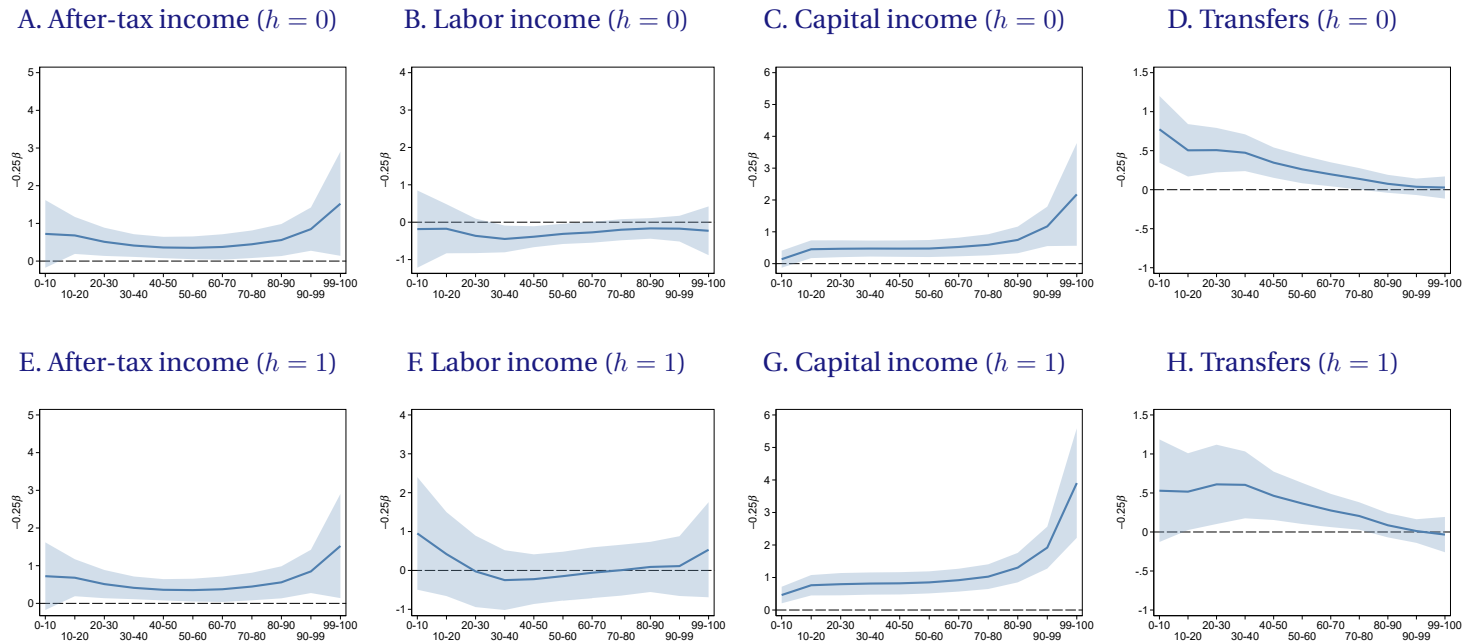
This figure shows the effects of a -25bp monetary policy shock on total after-tax incomes, as estimated using (1) on samples comprising only prime working-age individuals (26-65 years old; solid blue line) and all individuals older than 26 years (dashed green line), respectively. The estimation horizon is $h = 2$.

Figure C4: Labor-income results for continuously employed individuals ($h = 2$)



This figure shows the effects of a -25bp monetary policy shock on labor incomes, as estimated using (2) on the full sample (solid blue line) as well as on a sample comprising only continuously employed individuals (dashed green line). The latter comprises individuals whose labor income is at least half of the median of annual labor income in Sweden in both $t-1$ and $t+h$. The data on median labor income are from Statistics Sweden (2021). The estimation horizon is $h = 2$.

Figure C5: The effects of a -25bp shock on total after-tax income and its components ($h = 0$ and $h = 1$)



This figure shows the effects of a -25bp monetary policy shock on total after-tax income—as well as on the three main components of total pre-tax income—across the income distribution, as estimated using (1) and (2). The estimation horizon is $h = 0$ in Panels A–D and $h = 1$ in Panels E–H. Shaded areas represent 90 percent confidence bands.

Appendix D. Replication of Guvenen et al. (2017)

This appendix reports the results of a ‘scientific replication’ (Hamermesh, 2007) of the findings in Guvenen et al. (2017)—i.e., a re-examination of their findings using precisely the same econometric methods, but with data from a different institutional setting. The replication is based on the matched employer-employee database RAMS, compiled by Statistics Sweden based on administrative data collected from the Swedish Tax Authority.^{D1} RAMS is an annual panel comprising data on total labor income, main employer, and demographic characteristics for all residents in Sweden 16 years or older. The labor income reported in RAMS is the sum of earnings across all of an individual’s employers during a given year and includes wages, salaries, bonuses, stocks and exercised stock options, bonds, and taxable employee benefits. In keeping with the definition in Guvenen et al. (2017), self-employment income is excluded from the earnings measure. The outcome variable in all estimations is real earnings growth, defined as the log change in real earnings between years $t - 1$ and t . The nominal earnings figures in RAMS are deflated to real earnings using the GDP price deflator with 2010 as base year.

The sample covers the period 1987–2015 and is restricted to prime-age workers between 26 and 65 years old. In each year, the sample is sorted into four age groups (26–35, 36–45, 46–55, and 56–65) and twelve earnings bins (using cutoffs at percentiles 10, 20, ..., 90, 99, and 99.9). The sorting into earnings percentiles is based on past average earnings—defined as average annual real earnings over the years $t - 6$ to $t - 2$ —and is done conditional on gender and age group. For observations lacking earnings data in one or several years between $t - 6$ and $t - 2$, past earnings are calculated based on the longest consecutive period with available data, starting from year $t - 2$ and going backwards. The data required for computing earnings growth and past average earnings means that a worker needs to have positive earnings in at least years t , $t - 1$, and $t - 2$ to be included in the sample.

Workers’ exposure to systematic earnings risk are estimated in the form of “betas,” defined as the slope coefficients from regressions of real annual earnings growth on the two risk factors under consideration: real GDP growth and real stock returns. More specifically, the GDP beta for a worker belonging to a given gender-age-earnings group

^{D1}RAMS is one of the individual registries that goes into the construction of LISA, which the empirical analysis in the main part of the paper builds on.

g is estimated using the following regression specification:

$$\Delta y_{i,t} = \alpha_g + \beta_g \Delta y_t + \varepsilon_{i,t}, \quad (\text{D1})$$

where $\Delta y_{i,t}$ is the log real earnings growth of worker i from year $t - 1$ to t and Δy_t is the log real GDP growth from year $t - 1$ to t . The estimation of equation (D1) is carried out using pooled OLS, separately for each group g . Stock return betas are estimated using the same specification, but with real annual stock returns as regressor.^{D2}


Figure D1 plots GDP and stock return betas for 36–45 year old workers by gender, as well as for males by age group (the dotted lines represent 95-percent confidence intervals). Both GDP and stock return betas are U-shaped with respect to the earnings distribution, which is to say that workers in the top and bottom of the distribution are most exposed to aggregate earnings risk; this pattern holds for both males and females (although it is less pronounced for high-earning females), as well as within each age group for males. Throughout the earnings distribution, males and younger workers are more exposed to aggregate risk than females and older workers. The highest GDP beta, 3.81, is observed for 26–35 year old males in the lowest decile of the earnings distribution. This group also has the highest stock return beta together with 26–35 and 36–45 year old males in the top of the earnings distribution.

These cross-sectional patterns of earnings risk are qualitatively very similar to those for American workers reported by Guvenen et al. (2017). The *levels* of aggregate risk exposures are generally lower for Swedish workers than for American workers, however. For example, the GDP betas of 36-45 year old Swedish males in the bottom, middle, and top of the earnings distribution are 2.26, 0.40, and 1.90, respectively, whereas the figures for the corresponding groups of American workers are 2.88, 1.09, and 3.70 (i.e., about twice as high on average).

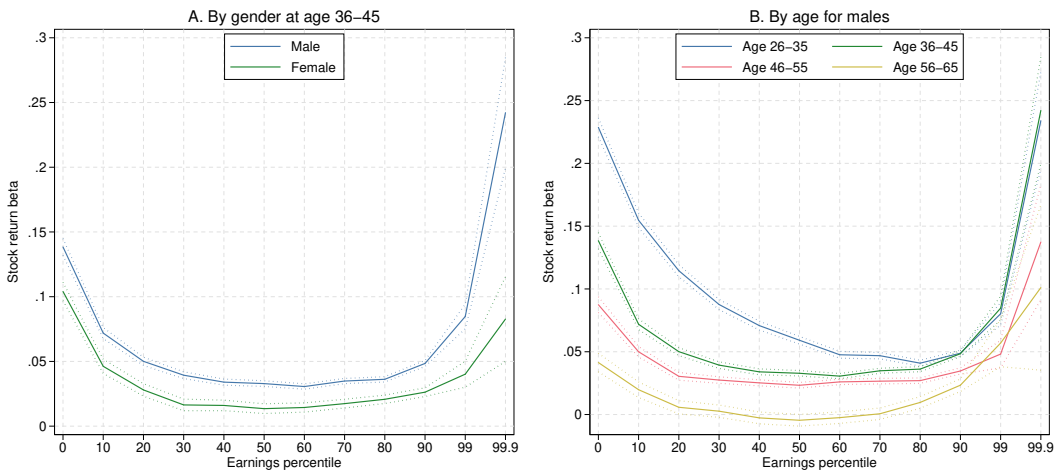
^{D2}Real annual stock returns are calculated based on the nominal Swedish stock return index compiled by Waldenström (2014), deflated by the GDP price deflator. Stock returns are aligned with earnings growth using the beginning-of-period convention, i.e., earnings growth from year $t - 1$ to t is aligned with real stock returns in year $t - 1$. This produces a correlation of real stock returns and real GDP growth of 0.70 over the period 1987–2015.

Figure D1: Worker betas

A. GDP betas by gender and age



B. Stock return betas by gender and age



This figure reports GDP betas (Panel A) and stock return betas (Panel B) by age, gender, and earnings group, as estimated using the specification in (D1). Dotted lines represent 95-percent confidence bands, computed using heteroskedasticity-robust standard errors.

Appendix E. Effects of Monetary Policy Shocks on the Values of Individual-Level Asset Holdings

In this section, we assess the effects of monetary policy shocks on the values of individual-level asset holdings. The direct approach for doing this would be to estimate our baseline specification with the growth in the value of asset holdings, expressed as a share of total-after tax income, as dependent variable. Our wealth data only spans the years 1999–2007, however, which is too short to obtain precise and plausible estimates using this approach.^{E1} Instead, we undertake an indirect, approximate assessment of the heterogeneous effects of monetary policy shocks on the value of asset holdings by combining our wealth micro data with estimates from the macroeconomic literature on the effects of monetary policy shocks on asset prices.

In our exercise, we consider two asset classes separately: stocks (including directly held stocks and equity mutual funds) and housing (including all types of real assets). For each asset class, we use our micro data to compute—for each of the eleven income groups—the average ratio of the holdings of the asset in question (at estimated market values) to after-tax income over the period 1999–2007. We then assume that the rate of return on each asset class is the same for all income groups in a given year, which allows us to take aggregate estimates of the effects of monetary policy shocks on asset returns from the empirical macroeconomic literature. Finally, by multiplying the estimates from the literature with the average ratios of asset holdings to after-tax income in the respective income groups, we obtain an estimate of the effect of monetary policy shocks on the values of asset holdings, expressed as a share of after-tax income, for each income group.^{E2} That is, we estimate the effect of a 25 basis points expansionary monetary policy shock on the value of the holdings of asset type j for an individual in income group g as

$$-0.25 \cdot \beta^j \cdot \frac{K_g^j}{Y_g^T}, \quad (\text{E1})$$

^{E1}Due to the existence of a wealth tax, the tax authorities collected information on individuals' holdings of financial and real assets from public registers and private institutions. When the wealth tax was abolished in 2007, the collection of wealth information stopped. For details on the wealth data, see, for example, Calvet, Campbell and Sodini (2007).

^{E2}Note that this approach only allows us to capture changes in asset holdings stemming from the effects that monetary policy shocks have on asset prices; changes stemming from the effects of monetary policy shocks on the composition of individuals' asset portfolios will, on the other hand, not be captured by our approach.

where $-0.25 \cdot \beta^j$ is the effect of a 25 basis points expansionary monetary policy shock on the price of asset type j , and K_g^j/Y_g^T is the average ratio of the holdings of asset type j to total after-tax income in income group g over the period 1999–2007. We take 1.5 percent as our estimate of the effect of a 25 basis points expansionary monetary policy shock on stock prices as well as on house prices.^{E3}

The results of this exercise are reported in Table E1. Panel A reports the average values of asset holdings by income group, Panel B the average ratios of asset values to total after-tax income, and Panel C the effects of a –25 basis points monetary policy shock on the values of asset holdings. First, note that while the values of housing and stock holdings both increase monotonously over the income distribution, the average ratio of assets to after-tax income is U-shaped over the income distribution for both asset types, as incomes increase somewhat faster than asset holdings over the bottom half of the distribution. The values of asset holdings then increase sharply in the top, both in absolute values and expressed as a share of total income: total asset values in the top percentile are on average 12 times as large in absolute value, and 2.5 times as large relative to after-tax income, as in the middle of the distribution. This is mainly due to the strong concentration of stock holdings in the top of the distribution; the value of stock holdings of individuals in the top percentile is on average 65 times as large as in the middle, or 14 times as large expressed as a share of after-tax income. Housing is, on the contrary, fairly equally distributed across the income distribution, at least when expressed relative to after-tax incomes.

Next, Panel C shows that an expansionary monetary policy shock leads to large increases in the values of asset holdings throughout the income distribution, but especially so in the very top. In the first nine deciles, a –25bp shock generates an increase in the value of total assets amounting to around five percent of after-tax income, almost entirely due to the increase in the value of housing. In the top percentile, the corresponding effect is more than 11 percent of after-tax income, in equal measure due to the increases in the values of housing and stock holdings. Interestingly, the effects of

^{E3}For the house-price estimate, we rely on Williams (2015), who reports the effect of a one percentage point contractionary monetary policy shock on real house prices to be –6 percent. Williams' own estimate is very close to the average estimate of –6.9 percent across the eleven studies included in his meta-analysis (Williams, 2015; Table 1). For the stock-price estimate, we take the average of Bernanke and Kuttner's (2005) main estimate of the effects of a 25bp expansionary shock on broad stock indices (1.2 percent) and the average of the four main estimates reported by Rigobon and Sack (2004) (1.7 percent); the average of the estimates from these two papers, then, amount to around 1.5 percent.

Table E1: The effect of monetary policy shocks on the value of asset holdings

	Income group										
	0-10	10-20	20-30	30-40	40-50	50-60	60-70	70-80	80-90	90-99	99-100
A. Average value of asset holdings (1000s)											
Total assets	448	516	541	571	612	669	755	884	1,137	1,841	7,979
- Housing	411	478	501	525	559	607	678	781	977	1,495	3,936
- Stocks	36	39	41	46	52	62	77	103	160	345	4,044
B. Ratio of average value of asset holdings to average after-tax income											
Total assets	3.70	3.27	3.08	2.98	2.94	2.98	3.11	3.33	3.80	4.65	7.56
- Housing	3.40	3.02	2.85	2.74	2.69	2.71	2.79	2.94	3.26	3.78	3.73
- Stocks	0.30	0.25	0.23	0.24	0.25	0.28	0.32	0.39	0.54	0.87	3.83
C. Effect of a –25bp monetary policy shock on the value of asset holdings											
Total assets	5.55	4.90	4.62	4.46	4.41	4.47	4.66	4.99	5.70	6.98	11.34
- Housing	5.10	4.54	4.27	4.11	4.04	4.06	4.19	4.41	4.90	5.67	5.60
- Stocks	0.45	0.37	0.35	0.36	0.38	0.41	0.47	0.58	0.80	1.31	5.75

Panel A reports the average values of total assets, housing, and stocks by income group in thousands of SEK (deflated to 2015 SEK using the GDP price deflator). Panel B shows the ratios of the average values of asset holdings to average after-tax income in each income group. Panel C reports the effects of a –25bp monetary policy shock on the values of asset holdings, as estimated using (E1). The numbers in Panels A and B are based on all individuals in the main estimation sample over the period 1999–2007.

monetary shocks on the value of stock holdings closely resemble the corresponding effects on realized capital gains reported in Table 2, both in level and in cross-sectional pattern. This is consistent with the observation that financial assets are traded—and thus that any gains or losses are realized—more frequently than housing assets are.