Five Facts about the Distributional Income Effects of Monetary Policy*

Niklas Amberg† Thomas Jansson‡ Mathias Klein§

Anna Rogantini Picco¶

May 1, 2021

Abstract

We use Swedish administrative individual-level data to document five facts about the distributional income effects of monetary policy. (i) The effects of monetary policy shocks are U-shaped with respect to 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 Klaus Adam, Martin Blomhoff Holm, Mikael Carlsson, Fabio Canova, Cristiano Cantore, Marco Del Negro, Michael Ehrmann, Veronica Guerrieri, Tor Jacobson, Stefan Laséen, Benjamin Moll, Ricardo Reis, Kasper Roszbach, Jiri Slacalek, Ulf Söderström, Oskar Tysklind, and Erik Öberg, as well as conference and seminar participants at the 11th ifo Conference on Macroeconomics and Survey Data, the Nordic Junior Macro Seminar Series, Norges Bank's 2020 Central Bank Macro Modelling Workshop, 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 individual-level income effects of monetary policy shocks. Taken together, these facts shed light on the overall distributional effects of monetary policy, as well as their underlying drivers.

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 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, except in the bottom decile. 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, entirely due to the *income composition channel*—that is, to the fact that capital income 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, quite stable over the income distribution.

We believe that these facts have relevance beyond Sweden for at least three reasons. 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. Secondly, the trends in income inequality in Sweden in recent decades are similar to those in most other developed economies, with large increases in the Gini coefficient as well as in top-income shares (see, e.g., Roine and Waldenström, 2015). Thirdly, we provide evidence that the cross-sectional heterogeneity in the unconditional aggregate earnings risk of Swedish workers is very 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 three dimensions relevant for the question at hand—namely, monetary policy, income inequality, and individuals' aggregate risk exposures—which speaks in favor of the external validity of our results.

Related literature. We contribute to the literature on the distributional income effects of monetary policy in several ways. In particular, most extant empirical evidence relies on survey data and/or time-series data on summary measures of income inequality, like the Gini coefficient (see, e.g., Coibion et al., 2017; Mumtaz and Theophilopoulou, 2017; Furceri, Loungani and Zdzienicka, 2018). Our findings suggest that such studies risk missing important heterogeneities in the effects of monetary policy: first, because survey data is typically top-coded and therefore cannot capture the substantial heterogeneity in the effects of monetary policy within the top of the income distribution; and second, because the large effects in the tails of the distribution often offset each other in summary measures of inequality and thus are difficult to identify in time-series settings. Indeed, our estimates imply that the Gini coefficient is virtually unaffected by monetary policy shocks, despite the pronounced heterogeneity at the individual level. Hence, 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: 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—motivated by Kaplan, Moll and Violante's (2018) HANK model—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.

Andersen et al. (2020) is more similar in terms of question and empirical approach, but there are also important differences. In particular, we provide more detail on the underlying drivers of the income effects of monetary policy by documenting the respective roles of the earnings heterogeneity and income composition channels. It is also worth noting that our results differ in several respects; most importantly, whereas we find that the income effects of monetary policy shocks are U-shaped, Andersen et al. find monotonically increasing effects over the income distribution. Exploring the extent to which institutional differences—for example, the fact that the Swedish central bank operates in a floating exchange-rate regime, whereas the Danish currency is pegged to the Euro—can account for the differing results would be interesting, but is beyond the scope of this paper.

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 incomes varies over the income distribution. We do this using the following econometric model:

$$\frac{Y^T_{i,t+h} - Y^T_{i,t-1}}{Y^T_{i,t-1}} = \sum_{g=1}^{11} G_{i,t,g} \cdot \left[\alpha^{T,h}_g + \beta^{T,h}_g \cdot \widehat{\Delta i}_t \right] + \varepsilon^{T,h}_{i,t}, \tag{1}$$

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 income Y^T_i 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; $\alpha^{T,h}_g$ is a group-specific intercept; and $\widehat{\Delta i}_t$ is the monetary policy shock in year t, which will be discussed in Section 2.3. $h = 0, 1, 2$ 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 $\beta^{T,h}_g$, which capture the percent change in total income over an h-year horizon for individuals in income group g, following a contractionary monetary shock of one percentage point.

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

$$\frac{Y^k_{i,t+h} - Y^k_{i,t-1}}{Y^T_{i,t-1}} = \sum_{g=1}^{11} G_{i,t,g} \cdot \left[\alpha_g^{k,h} + \beta_g^{k,h} \cdot \widehat{i}_t \right] + \varepsilon^{k,h}_{i,t}, \tag{2}$$

which is identical to (1), except that the numerator in the dependent variable is the change in one of the components k in total income between years $t-1$ and $t+h$, where $Y^T_{i,t} = \sum_k Y^k_{i,t}$. 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 incomes into the effects attributable to each component—i.e., we then have $\beta^T_{g,h} = \sum_k \beta^k_{g,h}$. Thus, the contribution of component k to the effect of monetary policy shocks on total incomes is given by $\beta^k_{g,h} / \beta^T_{g,h}$.

Thirdly, any heterogeneity in the effect of monetary policy on component k of total income is accounted for by some combination of (i) heterogeneity in the share of component k in total 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:

$$\frac{Y^k_{i,t+h} - Y^k_{i,t-1}}{Y^T_{i,t-1}} = \frac{Y^k_{i,t+h} - Y^k_{i,t-1}}{Y^k_{i,t-1}} \cdot \frac{Y^k_{i,t-1}}{Y^T_{i,t-1}}, \tag{3}$$

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 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, where each individual's actual income composition is replaced by the sample average. More specifically, we construct these variables by multiplying the original dependent variables by a group-level adjustment factor λ_g, defined as:

$$\lambda_g = \frac{Y^T_{g,t-1}}{Y^k_{g,t-1}} \cdot \frac{Y^k_{t-1}}{Y^T_{t-1}}, \tag{4}$$

where the first factor is the inverse of the average share of component k in total income in group g, and the second factor is the corresponding average share in the entire sample. Intuitively, the first factor approximately cancels out the individual's actual income share, while the second replaces it with the average income share in the sample.[1] This exercise thus 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.

2.2 Data, sample, and variable definitions

The analysis is based on administrative register data from LISA ("Longitudinal integrated database for health insurance and labour market studies"), 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—including those of the Tax Authority, the Public Employment Authority, and the Social Insurance Agency—and is thus, unlike self-reported survey data, virtually complete and free of measurement error.[2] 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, labor market indicators such as number of days spent in unemployment, 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 pre-tax income, defined as the sum of labor income, capital income, and transfers. 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. 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. All income variables used in the analysis are

[1] 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 $t+h$. An alternative approach would be to compute the percent change in income component k for each individual using the change in the inverse-hyperbolic sine of income, and then multiply the resulting growth rate with the respective sample income shares.

[2] For examples of recent papers using data from LISA, or from one or several of the individual registries that goes into the construction of LISA, see Akerman et al. (2013), Kolsrud et al. (2018), Saez, Schoefer and Seim (2019), and Busch et al. (2021).

deflated to real terms using the GDP price deflator with 2010 as base year.

We sort individuals into income groups g at an annual basis, using past average total income as a proxy for permanent income. More specifically, in each year t, individual i's permanent income is computed as her average total 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, which correspond to deciles of the distribution of past average 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 sample to prime-age individuals between 26 and 65 years old with positive total income, and the sample period to 1999-2018.[3] To limit the influence of outliers, we drop observations for which the growth in total income—or in one of its components—exceed 500 percent. The resulting final sample comprises 73.5 million individual-year observations and 6.4 million unique individuals. Descriptive statistics for the main income variables and demographic characteristics by income group are provided in Table B1 in Online Appendix B.

2.3 The monetary policy shock series

We construct our monetary policy shock series, $d\,i_t$, 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 re-

[3] We 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.

Figure 1: Construction of the monetary policy shock series

A. The poor-man's sign restriction

B. $i_m = \quad + \quad i_m^{TBill} \ {}^{0} + "_{i;t}$

C. Monthly shock series

D. Annual 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).

striction. [4] This restriction involves setting the monetary surprise to zero in cases where stock returns on announcement days move in the same direction as the surprise in the market interest rate. More speci cally, 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

[4] An alternative would have been to de ne monetary surprises based on STINA contracts—overnight index swaps denominated in SEK—but the data on these contracts only begin in 2003. Reassuringly, our shock series is closely correlated with an analogously constructed series based on STINA contracts for the period 2003–2018, as shown in Figure B1 in Online Appendix B. The choice of one-month T-bills as the basis of the monetary surprise series follows Flodén et al. (2020).

against the daily returns of the OMX Stockholm All Share Index on the days of monetary policy announcements.[5] Observations in the first and third quadrants, on the other hand, imply a positive comovement between interest rate and stock market changes, which runs counter to the conventional monetary policy transmission mechanism and thus suggests that the interest rate responses to these announcements are 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 $\Delta i_m^{TBill'}$:

$$\Delta i_m = \alpha + \beta \Delta i_m^{TBill'} + \varepsilon_{i;t}: \tag{5}$$

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.[6]

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 B.

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 incomes. The income groups are reported on the horizontal axis and the β^T, scaled

[5] On 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. As this is most likely due to measurement error, we use the two-day change in the T-bill rate for this annonucement date.

[6] In Online Appendix A, we show that an estimated proxy-VAR using our surprise series as an external instrument delivers impulse responses that are broadly in line with the textbook monetary policy transmission mechanism at the aggregate level.

by –0.25, on the vertical axis; shaded areas represent 90 percent confidence bands. While monetary policy shocks have large and statistically significant effects on total incomes across the entire income distribution, these effects are particularly large in the tails. More specifically, a 25 basis points reduction in the repo rate increases the total incomes of the poorest and richest individuals by 2.3 and 3.1 percent, respectively, over a two-year horizon, whereas the corresponding response for middle-income individuals is 0.6 percent. Hence, the effects of monetary shocks on total incomes are 4-5 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 almost three times 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 on each of the three main components in total income—labor income, capital income, and transfers. The response of labor income is large and statistically significant in the bottom two deciles, but small and statistically insignificant throughout the rest of the distribution. The capital-income response, on the other hand, is statistically significant across the entire income distribution, with the exception of the bottom decile. The effect is particularly large in the very top—for example, the capital income response is around seven 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 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 provide further details on the decomposition of the total-income effects in Section 3.3.

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 B2 in Online Appendix B shows, however, that very little changes if 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 Guvenen et al.'s (2017) results on the systematic earnings risk of workers using Swedish data. The results, reported in Online Appendix C, shows that the cross-sectional patterns of (un-

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

A. Total income

B. Labor income

C. Capital income

D. Transfers

This gure shows the effects of a –25bp shock on total income and its components across the income distribution for the estimation horizon $h = 2$, as estimated using (1) and (2). Shaded areas represent 90 percent con dence bands.

conditional) 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 conditional on monetary policy shocks 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 fol-

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

	Initial value	Two years after –25bp shock	Percent change
Gini coef cient	0.315	0.316	0.19
Top 1% income share	5.346	5.454	2.03
Top 10% income share	22.898	23.038	0.61
Standard deviation of log income	0.614	0.611	–0.41
Ratio of 90th to 10th percentile	4.001	3.985	–0.41
Ratio of 90th to 50th percentile	1.791	1.796	0.27
Ratio of 50th to 10th percentile	2.234	2.218	–0.68

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

lowing exercise. First, we compute the values of a number of commonly used measures of income inequality based on actual data for 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 total income in 2016 by $(1 \quad 0{:}25 \quad {}^{T}_{g;2})$, where g is given by the income group to which an individual belongs in 2016. [7] 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 for 2016.

The results are reported in Table 1. The Gini coef cient 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 over 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 small. On the other hand, the standard deviation of log income, as well as percentile

[7] Note that since the estimated ${}^{T}_{g;2}$ are negative—i.e., total incomes decline when the interest rate increases, and vice versa—the simulated income responses are positive.

ratios 90-10 and 50-10 all decrease, indicating that an expansionary monetary policy shock lowers income inequality.

These results imply that the most commonly used aggregate measures of income inequality—in particular the Gini coef cient—are not well-suited for characterizing the distributional effects of monetary policy. Fully understanding the distributional consequences of monetary policy instead requires looking at the impact of monetary policy over the entire income distribution, which can only be done with large-scale, uncensored individual-level administrative data like ours.

3.3 Decomposing the total-income effects

How does each income component 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 reports the ^{k}, as before scaled by –25bp, from the estimation of (2) with each of the three main components in total income, as well as their respective subcomponents, as dependent variables. The rows “Total income”, “Labor income”, “Capital income”, and “Transfer income” correspond to the coef cients plotted in Panels A–D of Figure 2.

To begin with, the leftmost column shows that of the 2.3 percent increase in total incomes for the poorest individuals following a –25 basis points shock, 2.0 percentage points is due to labor income, 0.2 to capital income, and 0.1 to transfer income. In the middle of the distribution, capital incomes account for about two thirds of the total-income response and transfers for the remainder, as the contribution of labor income is close to zero. For individuals in the top of the distribution, nally, the capital-income response accounts for 2.6 percentage points—or around four fths—of the total-income effect of 3.1 percent, with the remainder being due to labor income.

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 small average share in labor income over most of the income distribution (see Table B1

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
Labor income	2.0	0.9	0.3	0.0	0.0	0.1	0.2	0.2	0.2	0.2	0.6
- Wage income	2.0	0.8	0.3	0.0	0.0	0.1	0.2	0.2	0.2	0.2	0.6
- Self-employment income	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Capital income	0.2	0.3	0.3	0.3	0.3	0.4	0.4	0.5	0.7	1.1	2.6
- Realized capital gains	0.3	0.3	0.3	0.3	0.3	0.4	0.4	0.5	0.6	0.9	2.0
- Dividends and interest	–0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.2	0.6
- 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
Market income	2.3	1.2	0.6	0.3	0.3	0.4	0.6	0.7	0.8	1.3	3.2
Transfer income	0.1	0.3	0.4	0.4	0.3	0.2	0.1	0.1	0.0	0.0	–0.1
- Pensions	0.1	0.1	0.1	0.1	0.1	0.0	0.0	0.0	0.0	0.0	–0.1
- Unemployment income	–0.2	–0.2	–0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
- Other transfers	0.2	0.4	0.3	0.3	0.2	0.1	0.1	0.1	0.0	0.0	0.0
Total income	2.3	1.4	1.0	0.7	0.6	0.6	0.7	0.7	0.8	1.2	3.1

This table shows the contribution of each of the three main components in total income—as well as of their respective subcomponents—to the total income effects of a –25bp monetary policy shock, as estimated using (1) and (2).

in Online Appendix B). [8]

The capital-income response is decomposed into the parts due to realized capital gains, dividends and interest, and other capital income. Virtually the entire response is accounted for by realized capital gains in all income groups except for the top percentile, where dividends an interest account for 0.6 percentage points—or around one fourth—of the overall effect of 2.6 percent.

Finally, as the transfer-income response is statistically insigni cant across almost the entire income distribution, we refrain from drawing conclusions about the respec-

[8] Note, however, that self-employment income in the of cial Swedish income statistics only comprise income from self-proprietorships and trading partnerships; the incomes of individuals who are self-employed in incorporated rms are instead classi ed as wage income or dividends. Hence, our data likely understates the role of self-employment income.

tive contributions of its subcomponents. The estimated effects for the subcomponents are nevertheless included in Table 2 for completeness.

3.4 Income composition versus within-component heterogeneity

Heterogeneity in the effects of monetary policy shocks on component k of total income is, as discussed in Section 2.1, accounted for by some combination of heterogeneity in the share of component k in total 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 total income varies over the income distribution. This is done in Panel A of Figure 3. Labor income constitutes the largest share of total income in all income groups 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 capital income is small and roughly constant over most of the distribution, but increases sharply in the top; for example, the capital-income share is almost ten times as large in the top percentile as in the middle of the distribution (18 percent versus two percent). The share of transfer income, finally, is large in the bottom (46 percent of total income in the first decile), but then decreases monotonically over the income distribution.

Now, turning to the formal analysis, we compare the actual estimates of ^{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 already reported in Figure 2, while the dashed green lines are the counterfactual estimates described in Section 2.1.

Figure 3: Income composition versus within-component heterogeneity

A. Income composition by group

B. Labor income

C. Capital income

D. Transfers

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 shocks—namely, strong effects in the bottom of the distribution, but small and insignificant effects in the middle and top—is entirely accounted for by the earnings heterogeneity channel. That is, the heterogeneity is due to the fact that labor incomes are much more sensitive to monetary shocks in the bottom of the income distribution than in the middle and top. The income composition channel, on the other hand, strongly attenuates the heterogeneity in the labor-income response.

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 seven times larger than in the middle—the counterfactual estimates are almost constant over the income distribution, with the exception of the bottom decile. The difference between the actual and counterfactual estimates are striking in the top: in the top percentile, for example, the counterfactual estimate is more than ve times as small as the actual estimate. This implies that the heterogeneity in the capital-income response is entirely due to the income composition channel—that is, to the fact that (i) capital income responds particularly strongly to monetary shocks, and (ii) that capital income constitutes a larger share of total income in the top of the distribution. The sensitivity of capital incomes to monetary shocks is, on the contrary, the same across virtually the entire 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 distribution—such as the savings redistribution, nancial segmentation, and portfolio channels (Coibion et al., 2017)—are not quantitatively important.

Panel D, nally, 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 insigni cant and of small magnitudes. In sum, the heterogeneity 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. In particular, we have shown that the effects of monetary policy shocks on individuals' total 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.

The facts presented in the paper are directly relevant for the policy debate on the distributional effects of monetary policy that has taken place in recent years. They should also be useful for the macroeconomic literature studying monetary policy in HANK models, since they can be used as targets to differentiate between competing classes of models. For example, the importance of labor-market outcomes for low-income individuals might ask for incorporating search and matching frictions into HANK models and allowing the degree of these frictions to vary across the income distribution.

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 new insights on the distributional effects of, for example, scal 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.

Akerman, Anders, Elhanan Helpman, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen Redding. 2013. “Sources of Wage Inequality.” *American Economic Review*, 103(3): 214–19.

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/.

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.

Busch, Christopher, David Domeij, Fatih Guvenen, and Rocio Madera. 2021. “Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance.” *American Economic Journal: Macroeconomics*, forthcoming.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2017. “Innocent Bystanders? Monetary policy and inequality.” *Journal of Monetary Economics*, 88(C): 70–89.

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.

Flodén, Martin, Matilda Kilström, Jósef Sigurdsson, and Roine Vestman. 2020. “Household Debt and Monetary Policy: Revealing the Cash-Flow Channel.” *Economic Journal*, fortcoming.

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.

Guvenen, 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.

Kolsrud, Jonas, Camille Landais, Peter Nilsson, and Johannes Spinnewijn. 2018. “The Optimal Timing of Unemployment Bene ts: Theory and Evidence from Sweden.” American Economic Review, 108(4-5): 985–1033.

Laséen, Stefan. 2020. “Monetary Policy Surprises, Central Bank Information Shocks, and Economic Activity in a Small Open Economy.” Sveriges Riksbank Working Paper No. 396.

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.

Roine, Jesper, and Daniel Waldenström. 2015. "Long-Run Trends in the Distribution of Income and Wealth." In *Handbook of Income Distribution*. Vol. 2, 469–592.

Saez, Emmanuel, Benjamin Schoefer, and David Seim. 2019. "Payroll Taxes, Firm Behavior, and Rent Sharing: Evidence from a Young Workers' Tax Cut in Sweden." *American Economic Review*, 109(5): 1717–63.

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."

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.

Online Appendix for "The Heterogeneous Income Effects of Monetary Policy"

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

Appendix A. Aggregate Effects of the Monetary Policy Shock

To justify our choice of the monetary policy shocks, we estimate a proxy-VAR to study the induced aggregate dynamics. In particular, we use the monthly monetary policy surprise series as 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 as 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 takes the value of 9:02 such that weak instrument problems are unlikely to be a major concern for our 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 shaded areas indicate 90% confidence bands obtained from 1,000 bootstrap repetitions, and the darker shaded areas indicate 68% confidence bands. 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 supports our shock construction for studying the effects of monetary policy at the individual level.

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. Shaded areas indicate 68% and 90% bootstrapped con dence intervals. The unit of the horizontal axis is a month and the sample is 1999M1-2018M12.

Appendix B. Additional Tables and Figures

This appendix provides additional tables and gures referred to in the main text of the paper. Table B1 presents descriptive statistics by income group; Figure B1 a comparison of our monetary shock series with an analogously constructed series based on STINA contracts; Figure B2 the total-income responses when the nancial crisis is excluded from the estimation sample; and Figure B3 the effects of a –25bp monetary policy shock on total income and its components for estimation horizons $h = 0$ and $h = 1$.

Table B1: 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 income (1,000s)	152	207	238	266	293	321	353	395	466	669	1,718
B. Average shares of total income											
Labor income	0.48	0.58	0.68	0.77	0.83	0.86	0.88	0.89	0.89	0.88	0.76
- Wage income	0.41	0.53	0.65	0.74	0.80	0.84	0.86	0.87	0.87	0.86	0.74
- Self-employment income	0.07	0.05	0.03	0.03	0.02	0.02	0.02	0.02	0.02	0.02	0.02
Capital income	0.05	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.04	0.07	0.18
- Realized capital gains	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.05
- Dividends and interest	0.04	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.04	0.12
- Other capital income	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
Transfer income	0.46	0.40	0.30	0.21	0.15	0.12	0.09	0.08	0.07	0.05	0.06
- Pension income	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.04
- Unemployment bene ts	0.06	0.09	0.08	0.05	0.03	0.02	0.02	0.01	0.01	0.01	0.01
- Other transfer income	0.32	0.26	0.19	0.14	0.10	0.08	0.06	0.05	0.04	0.03	0.01
C. Other characteristics (means)											
Male	0.42	0.34	0.31	0.34	0.41	0.51	0.60	0.65	0.69	0.75	0.84
Age	40.4	41.0	42.1	43.1	44.1	44.7	45.2	46.0	46.8	48.4	50.6
Less than high-school education	0.26	0.18	0.16	0.15	0.14	0.15	0.15	0.13	0.10	0.06	0.04
High-school education	0.40	0.48	0.56	0.57	0.57	0.55	0.52	0.48	0.41	0.29	0.21
Post-secondary education	0.34	0.33	0.29	0.28	0.28	0.31	0.34	0.39	0.49	0.66	0.74
Unemployment days/year	22.4	19.5	17.7	13.9	10.2	8.1	6.4	5.0	4.2	3.8	3.3

Figure B1: Comparison of monetary policy shock series (1M T-bill versus STINA)

This gure compares the monetary policy shock series used in the paper (solid blue line) with an analogously constructed shock series based on STINA contracts (dashed green line). The data on STINA surprises are from Laséen (2020).

Figure B2: Total-income results with and without nancial crisis ($h = 2$)

This gure shows the effects of a –25bp shock on total incomes, as estimated using (1), when the nancial crisis is included and excluded, respectively. The estimation horizon is $h = 2$.

Figure B3: The effects of a –25bp shock on total income and its components ($h = 0$ and $h = 1$)

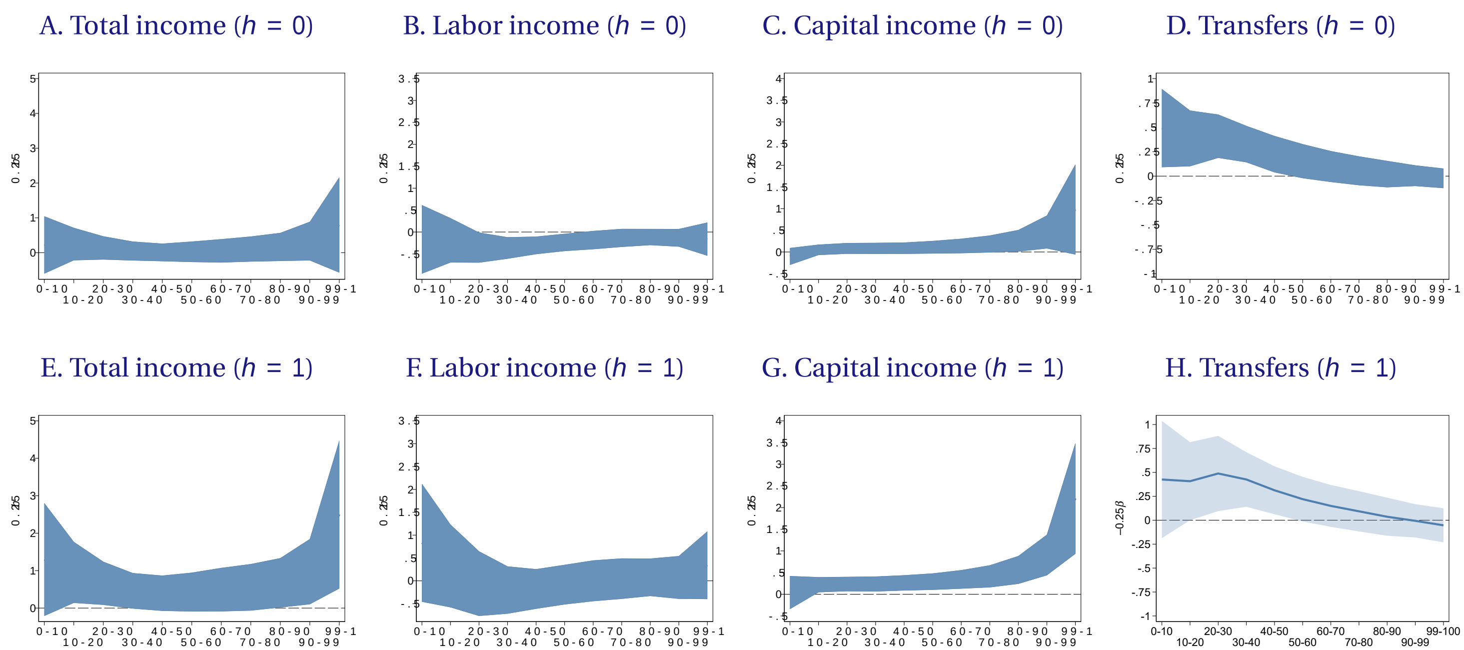

This figure shows the effects of a –25bp shock on total income and its components across the income distribution for the estimation horizons $h = 0$ and $h = 1$, as estimated using (1) and (2). Shaded areas represent 90 percent confidence bands.

Appendix C. 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.[C1] 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

[C1] 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.

β_g is estimated using the following regression specification:

$$\Delta y_{i,t} = \alpha_g + \beta_g \Delta y_t + \varepsilon_{i,t}, \tag{C1}$$

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 (C1) 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.[C2]

Figure C1 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).

[C2] 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.