Monetary policy shocks and inflation inequality

Christoph Lauper* Giacomo Mangiante†

November 2021

Abstract

We evaluate household-level inflation rates since 1980, for which we compute various dispersion measures, and we assess their reaction to monetary policy shocks. We find that (i) contractionary monetary policy significantly and persistently decreases inflation dispersion in the economy, and that (ii) different demographic groups are heterogeneously affected by monetary policy. Due to different consumption bundles, middle-income households experience higher median inflation rates, which at the same time are more reactive to a contractionary monetary policy shock, leading to an overall convergence of inflation rates between income groups. These results imply that (iii) the impact of monetary policy shocks on expenditure inequality is significantly more muted once we control for differences in individual inflation rates.

Keywords: monetary policy, inflation inequality, redistributional effects
JEL classification: E31, E52

*University of Lausanne, Lausanne, Switzerland, email: christoph.lauper@unil.ch
†University of Lausanne, Lausanne, Switzerland, email: giacomo.mangiante@unil.ch
We would like to thank Florin Bilbiie, Jean-Paul Renne, and Andreas Tischbirek for their valuable comments and guidance. We are also grateful to Philippe Bacchetta, Christina Patterson, Kenza Benhima, and fellow UNIL Ph.D. students for helpful suggestions that significantly improved the paper. Finally, we would like to thank the participants of the conferences organized by the International Association for Applied Econometrics, the Swiss Society of Economics and Statistic, the Society of Economics of the Household, the Ph.D. - Economics Virtual Seminar, the VMACS (Virtual Macro Seminar), the Lausanne Macro Reading Group, the 9th Italian Congress of Econometrics and Empirical Economics, the 61st Annual Conference of the Italian Economic Association and the 15th Economics Graduate Students Conference for precious feedback.

1 Introduction

Consumer heterogeneity has become increasingly important in macroeconomic research, both from a theoretical and empirical point of view. This can be seen in the rise of Heterogeneous Agent New Keynesian models (HANK, see Kaplan et al., 2018) which show that differences among households have important repercussions on the transmission of monetary policy. At the same time, Coibion et al. (2017) document that income and expenditure inequality increase after a contractionary monetary policy shock. On a more general level, these results show that the focus in monetary economics has shifted considerably from aggregate to individual effects.

Any individual variable needs to be deflated by a measure of inflation to get meaningful real variables. However, almost all research on consumer inequality uses the aggregate inflation rate to deflate the household-level variables, thereby abstracting from important differences in individual inflation rates. We try to fill this gap. Analyzing the differences and dynamics of inflation heterogeneity is essential for a thorough and unbiased view on consumer heterogeneity.

In this paper, we document that differences in individual inflation rates are sizable and that they are systematically different for different demographic groups. We show that monetary policy plays a significant role in the evolution of the cross-sectional inflation distribution: following a contractionary monetary shock, the dispersion and skewness of the distribution decrease.

Moreover, we find that the level and volatility of the individual inflation rates strongly depend on household expenditure, salary, and income. This non-random behavior of inflation rates has direct effects on monetary policy. By defining inflation inequality as the cross-sectional dispersion of individual inflation rates along the expenditure, salary or income distribution, we demonstrate that contractionary

monetary policy shocks lead to a substantial and persistent decrease of inflation inequality in the medium term.

To obtain individual inflation rates, we exploit the fact that consumers purchase different consumption bundles, with different shares spent on different subcategories of the CPI.[1]

Assuming all households face the same prices for the same categories of goods and services, we use item-level price data from the Bureau of Labor Statistics (BLS). We combine it with individual expenditure data from the Consumer Expenditure Survey (CEX) to create inflation rates at the household level. This allows us to study the distribution of individual inflation rates across a representative sample of households that includes the same range of goods and services that is used in computing the Consumer Price Index for the U.S. We analyze the dynamic responses of different inflation dispersion measures to the exogenous monetary policy shocks derived by Romer and Romer (2004) and extended by Coibion et al. (2017).

We find that a contractionary shock to monetary policy significantly and persistently decreases inflation dispersion. This finding is robust across many different reference periods and empirical specifications. To assess the asymmetry of the response, we repeat the analysis on different skewness measures. The distribution's right tail converges more strongly than the left tail toward the median inflation rate, implying that high-inflation households see their inflation rates be more reactive to monetary shocks. Moreover, household expenditure, salary, and income are strongly correlated with the response of individual inflation. Inflation of low and middle-income households is more reactive to shocks than that of high-income households and therefore decreases more strongly after a contractionary monetary policy shock. At the same time, the latter group's median inflation rate tends to be lower. This

[1]The categories are 21 fairly broad baskets of goods and services, for which we can match the expenditure data to the price data (see chapter 2).

results in a decline in inflation inequality across income groups. The same results hold for salary and expenditure deciles, confirming the role of endowments on individual inflation rates. Finally, we show that neglecting inflation heterogeneity leads to biased estimates of the impact of monetary policy on expenditure inequality.

Related literature

This paper contributes to two strands of the literature. First, our results complement the large body of empirical evidence on the relationship between monetary policy and inequality. With an approach analogous to the one we adopt, Coibion et al. (2017) demonstrate for the U.S. how consumption and income inequality increase following a contractionary shock. Similar findings have been confirmed in other countries as well, and in different time periods (e.g. Mumtaza and Theophilopoulou, 2017, for the United Kingdom and Samarina and Nguyen, 2019 for the Euro Area).

The second strand is the growing literature on the heterogeneous responses of households to monetary policy shocks across demographic characteristics. An active part of the research community has focused its attention on expenditure inequality. Using CEX data as well, Wong (2019) documents that young people adjust their consumption more than middle-aged and older households. Exploiting differences in the housing tenure of the survey respondents, Cloyne et al. (2019) show that households with mortgage debt are the most sensitive group to shocks whereas the consumption of homeowners without debt is basically unaffected by the change in the interest rate.

Less attention has been paid to heterogeneity in terms of inflation rates across demographic groups. Previous studies include, among others, Johannsen (2014) (using CEX data) and Kaplan and Schulhofer-Wohl (2017) (using Nielsen scanner data) who document substantial cross-sectional dispersion in household inflation

rates. Particularly related to our results, Jaravel (2019) found for the period 2004-2014 that high-income households are exposed to much lower inflation rates than low and middle-income households.

Finally, Cravino et al. (2020) show how following a contractionary shock high-income households' inflation rate reacts significantly less than the one of middle-income households. Our paper differs over different dimensions: first of all, we compute the inflation rates at the household level instead of at the percentile level. This allows us to provide new evidence on the role played by monetary policy on the cross-sectional inflation distribution. Furthermore, when studying the monetary policy transmission across demographic groups, our focus is not on the absolute value of the impulse responses but rather on the relative response among groups. Conditioning on the group median inflation rate over time, we are then able to document whether inflation dispersion increases or decreases in response to a monetary shock.

The paper is structured as follows. Section 2 describes the dataset used, as well as the construction of individual inflation rates and dispersion measures. In Section 3 we discuss the empirical strategy and show the main results in terms of the impact of monetary policy shocks on the cross-sectional inflation distribution. Section 4 studies the heterogeneous responses across different demographic groups. In section 5, we perform a battery of different robustness checks to evaluate the reliability of our findings. Section 6 concludes.

2 Individual inflation rates

In this section, we compute individual inflation rates at the household level. While Cravino et al. (2020) and others have calculated inflation rates at the percentile level, focusing on the differences between different demographic groups, and therefore

abstracting from within-group dispersion, we want in a first step to focus on the heterogeneity across all households.

We exploit the differences in consumption patterns between different households and apply good-level price indices to expenditure categories in order to retrieve individual, household-level inflation rates.

There are three steps needed for the computation of any inflation rate. We need information on prices for different goods. Then, we need detailed information on (individual) consumer expenditures, which allows computing the share of different goods in an aggregated index and therefore providing weights. [2] Thirdly, statistical agencies have to decide on a methodology to combine price data to get a meaningful measure of inflation. In the following, we discuss each step separately.

2.1 Inflation data

We use data from the Consumer Price Index (CPI) as computed by the BLS at a monthly frequency. In particular, we use the not-seasonally-adjusted *US City Average for all urban consumers* (CPI-U). The BLS collects price data on 211 different subgroups of goods and services, which they call item strata. This is the most disaggregated level for which it publishes information on prices. However, these item strata over the period from 1980 to today undergo regular revisions or their definition is changed. Some disappear entirely and some get newly introduced since new goods get introduced and old ones disappear. For this reason and for data availability, we need to combine these basic price indices to more aggregate ones. We follow Hobijn and Lagakos (2005) in creating 21 indices, for which we get consistent inflation rates during our time sample. We will call these inflation rates of subgroups of the consumer basket *inflation subindices*[3]. The construction of these

[2] The CEX proves rich enough to provide data on expenditures, going back to 1980

[3] The list and definitions of these subindices can be found in Appendix A.1.

inflation rates is subject to a tradeoff between consistent and sufficiently long time series on one hand, and finely disaggregated time series that catch as much of the difference in inflation as possible.

In Table 1 we report the mean, median, standard deviation, the 10th and the 90th percentile of the 21 inflation subindices we compute, as well as of the Official CPI-U. The observed sectoral inflation heterogeneity will be one of the key components in explaining the evolution of inflation dispersion. Households spend different shares of their overall expenditure on each category and, since these categories differ in terms of price volatility and price level, this will lead to differences in terms of experienced inflation.

The biggest limitation of using inflation subindices is that they are not individual prices. While we capture the inflation that is due to different consumption baskets, we are not able to capture inflation differences within a subindex. It is conceivable that taking the category *Food away* as an example, that high-end restaurants have different price developments than low-end ones.

This problem is circumvented with Nielsen scanner data. The dataset reports product-level information on both prices and quantities so it is more granular than the CEX data. However, two major limitations made the Nielsen data not a viable solution for our analysis. First of all, the data covers only purchases in department stores, grocery stores, drug stores, convenience stores, and other similar retail outlets which account for approximately 15% of the total household expenditures. Moreover, the dataset is available only from 2004 onward.

In what follows, we have to find reliable weights with which we can combine the inflation subindices to get household-level inflation rates across all items.

2.2 Expenditure data

For the computation of expenditure weights, we use the CEX provided by the BLS. This is the same dataset that is used to compute the official CPI of the U.S. The CEX is a quarterly survey of household expenditures and is divided into a diary and an interview survey. The diary survey covers small expenditures on daily items over a period of two weeks. The interview survey is more comprehensive, with detailed questioning every three months yielding up to a year of data for a single household. Since our goal is to get inflation rates that are as comprehensive as possible, we will solely rely on data from the interview survey.

There are some limitations to the CEX data. The BLS removes consumption data from the 100th percentile (it is top-coded) to ensure anonymity. Also, since we deal with survey data, there are likely more measurement errors in the CEX compared to other data sources.[4] However, the CEX allows us to get a comprehensive picture of virtually all consumer expenditures and is also sufficiently large in the time dimension (starts in 1980) and along the cross-section (roughly 5000-7000 households each month).

Like the inflation subindices, we aggregate the expenditure data into 21 groups,[5] matching the classification of the CEX with the one from the price indices. In the next step, we aggregate the monthly expenses to yearly expenses. By doing this, we get rid of seasonal patterns in expenditures, while at the same time "averaging out" extraordinary expenses and hence improving the quality of our data. With this approach, most variation in individual inflation rates comes from price changes, rather than from changes in consumption patterns. Hence, the variation in individual inflation rates is mainly driven by the dynamics of sectoral inflation rates, as opposed

[4] cf: Bee et al. (2013) for an assessment of the quality of our consumer dataset.

[5] We have to alter the *Housing* group and omit the *Vehicle* group altogether for reasons specific to the group. See the Appendix for details.

to being driven by changes in the consumption bundle, as we intend.[6] Since we use inflation data that is the same for all households, it is important to stress the central importance of the variation in individual expenditure shares. All the dispersion in individual inflation rates is due to the variation in expenditure shares. Luckily, expenditure shares show large variation that can be explained to a large part by differences in total expenditure, income, or salaries. In order to show the variance in expenditure shares, we group the households in deciles of total expenditure and employ a correspondence analysis to display the differences in expenditure shares for all categories.[7]

2.3 Computation of individual inflation rates

In a third step, we combine the expenditure data with the inflation data. For this, we compute consumption shares w_j^i for household i and item subgroup j, which are calculated by dividing the yearly consumption expenditure in a certain period by the total expenditure reported in the same period. In the baseline analysis, we use all 21 categories. To assess the robustness of our results, we have various scenarios that exclude parts of the consumption basket (see Section 5). Then, we compute the individual inflation rate for household i as:

$$\pi_{t-k,t}^i = \sum_{j \in J} w_j^i \pi_{j,t-k,t} \tag{1}$$

where j denotes the item subgroup as defined in section 2.2. The inflation rate of the subindex for good j in period t with base period $t-k$ is denoted by $\pi_{j,t,t-k}$. We

[6]The dataset allows us to exploit differences in inflation due to different consumption bundles. However, we avoid exploiting the short-term dynamics in expenditure shares, since we expect the quality of the month-to-month data to not be sufficiently good. Households are only interviewed for a maximum of 12 months, which does not allow us to capture substitution effects.

[7]In the Appendix, table 2 displays the actual weights for individual sectors along expenditure, income and salary deciles.

Figure 1: Correspondence analysis of the variation in weights for different expenditure deciles

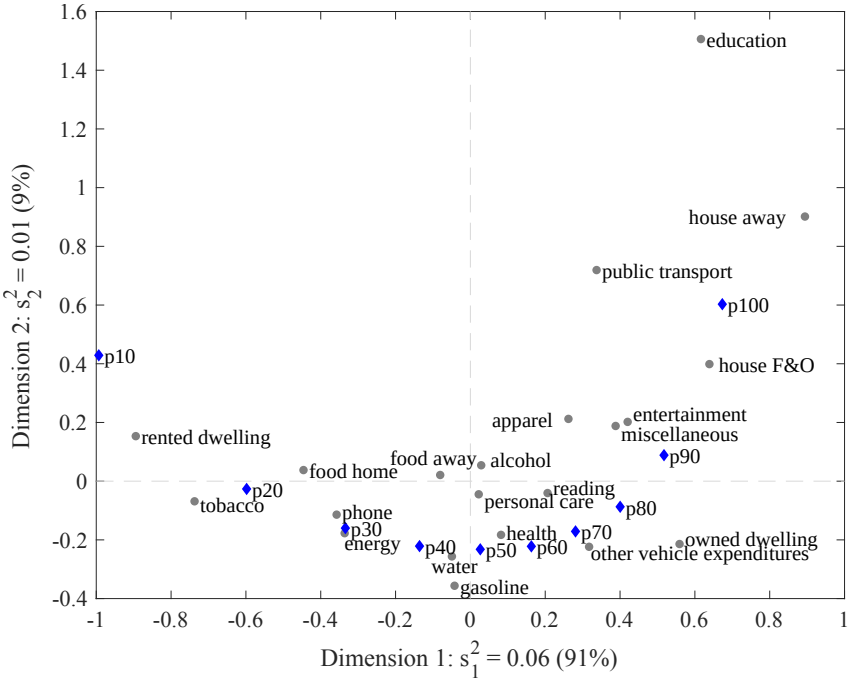

Notes: The correspondence analysis displays the scores for the two largest principal components of the weights for all 21 sectors used in our analysis. The first principal component, on the horizontal axis, accounts for 91% of the variation. It can be interpreted as showing linear differences between expenditure deciles. That is, goods on the left-hand side of the origin are purchased to a higher share by low-expenditure households, and vice versa. The vertical axis explains 9% of the variance. Sectors at the top of the figure (with a high y-value) are purchased relatively more by the poles of the distribution, whereas goods at the bottom (negative y-value)of the figure are purchased more by the middle of the distribution (e.g. gasoline) The distribution of weights is very similar across income deciles.

set $k = 12$, meaning year-on-year inflation rates, which removes seasonality in the inflation subindices. Additionally, we winsorize the individual inflation rates at the 1st and the 99th percentile. In the next step, we analyze the statistical properties of individual inflation rates.

2.4 Properties of individual inflation rates

We assess the validity of the measures of individual inflation computed above by comparing the official CPI inflation rate with the median of individual inflation rates. The scatter plot of the calculated household-specific rates of inflation depicts the dispersion of individual inflation rates (Figure 2).[8]

Figure 2: Official CPI inflation, dispersion, and median of individual inflation rates

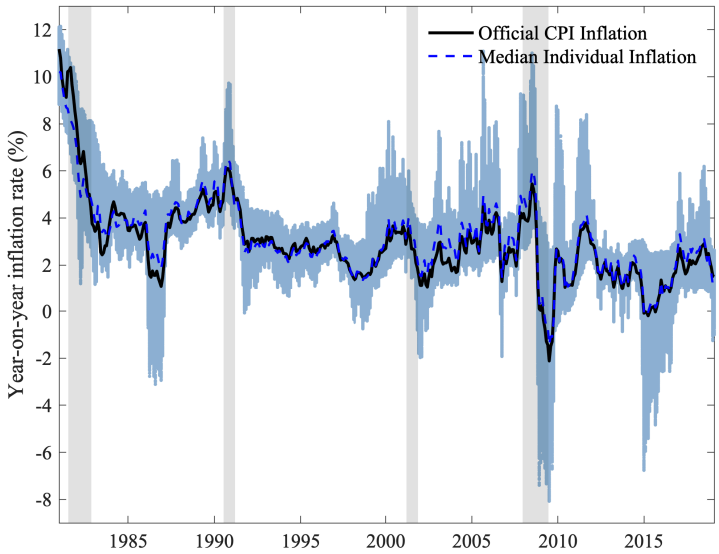

Notes: The scatter plot and median individual inflation rate are computed using winsorized data, meaning that the top and bottom 1% of household-level inflation rates at every point in time are excluded. The gray shaded areas depict U.S. recessions.

On the one hand, the median of the distribution of household-specific rates of inflation closely tracks the headline value of CPI inflation. Hence, our approach gives, in an aggregate world, very similar results to the official CPI inflation rate. This result shows why for many years economic models mainly focused on the representative agent: indeed, the time series of the experienced inflation for the

[8] Similar results are obtained for the mean of the distribution.

"median household" can be considered a quite good approximation of the aggregate economy.

On the other hand, the scatter plot in the same figure reveals how much information is lost when ignoring the heterogeneity across households. Not surprisingly, macroeconomic models have been expanded to include heterogeneity in consumption, wages, asset portfolio composition, and many more. However, most models still abstract from inflation differences and implicitly assume that households are exposed to the same inflation rate. Figure 2 seems to strongly reject this assumption.

2.5 Measures of dispersion

To evaluate how monetary policy shocks affect inflation dispersion in the United States, we construct three different measures of dispersion: the cross-sectional standard deviation, the difference between the 90th percentile and the 10th percentile (depicted as 90th-10th, henceforth), and the cross-sectional interquartile range (IQR) for the period 1980M1-2007M12. Our series deliberately stops before the recent financial crisis to exclude the zero lower bound period.

To avoid that the change in the survey composition might affect our results, we calculate the variation in the inflation dispersion measures on the households present in both periods. Therefore, when we calculate the change in cross-sectional standard deviation from t to $t+1$, we do it only for the households who are present during both periods. Sampling weights are applied throughout the analysis.

The top plot in Figure 3 shows the historical evolution of the three measures of dispersion, together with U.S. recessions. The three variables are highly correlated, suggesting that a normal distribution approximates the computed individual inflation rates very well. Despite using a different time period and alternative CPI categories, the time series are comparable in magnitude to what Johannsen (2014) found. As

Figure 3: Historical series of inflation dispersion measures and overall skewness

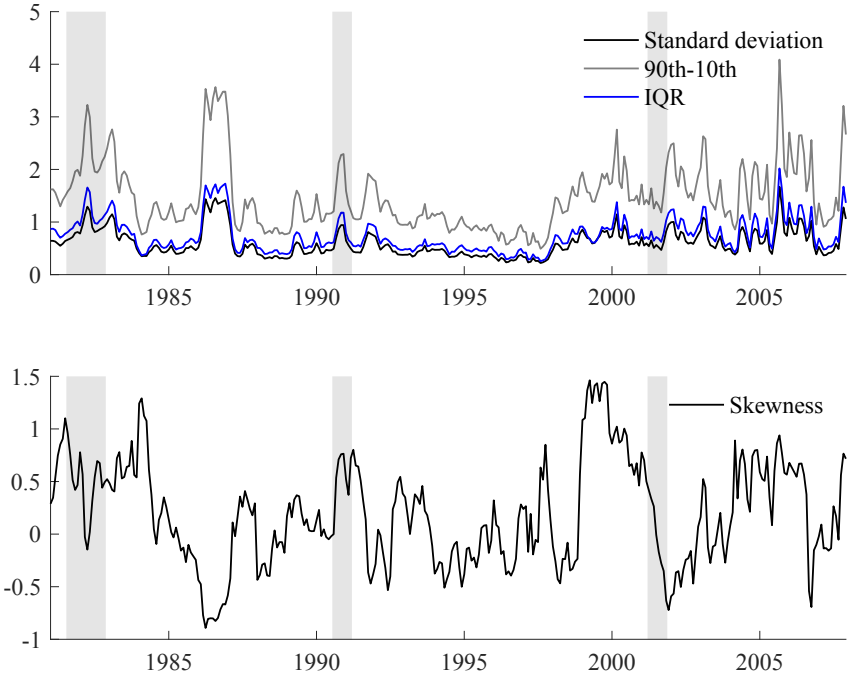

Notes: In the top plot we show the evolution of inflation dispersion measured using the cross-sectional standard deviation, the difference between the 90th and the 10th percentile of the cross-sectional distribution, and the IQR. The bottom plot reports the time series for overall skewness of the inflation distribution. All the series refer to the period 1981M1:2007M12. The gray shaded areas depict U.S. recessions.

one can notice, inflation dispersion increases during U.S. recessions of the early 80s and 90s suggesting some sort of correlation with the business cycle in the economy.

The bottom plot displays the time series for the overall skewness of the cross-sectional inflation distribution. The skewness and the dispersion measures are not strongly correlated. Therefore, the study of both second and third moments will convey complementary information regarding the impact of contractionary monetary policy shock on the shape of the distribution[9].

[9]See Appendix B for the contribution of the energy sector on the evolution of the dispersion measures.

3 The impact of monetary policy shocks on inflation dispersion

In this section, we present the results of our empirical analysis. We first study whether and to what extent monetary policy shocks influence aggregate inflation dispersion. We then investigate more in-depth how the inflation distribution reacts to contractionary shocks by focusing on the distance between the two tails and the median as well as on the overall skewness.

3.1 Methodology

In the baseline specification, we adopt the Local Projection (LP) method developed by Jordà (2005). In particular, we estimate a series of regressions of the dependent variable over different horizons on the monetary policy shock in period *t* and controlling for the lags of the shock as well as of the dependent variable as in Coibion et al. (2017) and Cravino et al. (2020):

$$x_{t+h} - x_{t+h-1} = c_h + \beta_h e_t^{RR} + \sum_{j=1}^{J} \theta_{h,j}(x_{t-j} - x_{t-j-1}) + \sum_{i=1}^{I} \gamma_{h,i} e_{t-i} + \epsilon_{t+h} \quad (2)$$

where x is the variable of interest. The monetary policy shocks are denoted by e_t^{RR} and c_h is a vector of horizon dummies. In line with the literature, we include as control 48 lags of the shocks and 6 lags of the dependent variable. The coefficient β_h for $h = 1, ..., H$ gives the response of the dependent variable at time *t*+*h* to a monetary policy shock at time *t* and is used to generate the accumulated impulse response to a 1 percentage point contractionary monetary policy shock.

To identify unanticipated changes in the short-term interest rate we use the monetary policy shock series devised by Romer and Romer (2004, henceforth called

R&R shocks), and extended by Coibion et al. (2017). The shock series covers the survey sample periods from 1980 to 2007 [10].

3.2 Analysis

To evaluate the overall effects of a contractionary monetary policy shock on inflation dispersion, we estimate equation (2) using the cross-sectional standard deviation as the baseline measure of inflation dispersion[11]. The impulse responses are computed over a horizon of 48 months using data from 1980M1 to 2007M12 and standard errors are corrected as in Driscoll and Kraay (1998) to allow for arbitrary serial and cross-sectional correlation across horizons and time. For each impulse response, we present one and 1.65 standard deviation confidence intervals.

The results are reported in Figure 4. The top panel shows the responses of the annual inflation rate computed by the BLS (black line) as well as of the median inflation rate across households: following a contractionary shock, the annual rate decreases by approximately 1.5 percentage points, a magnitude in line with the literature. As one might have expected looking at Figure 2, the response of the median inflation rate closely matched the response of aggregate inflation.

In the middle panel, we show the impulse response of our dispersion measure: dispersion decreases after a contractionary monetary policy shock and remains persistently below zero. Looking at the one and 1.65 standard deviation confidence intervals we can easily reject the null hypothesis that the coefficients are equal to zero for the horizon considered. Therefore, the impulse response strongly suggests

[10] Coibion (2012) shows how the Romer and Romer (2004) approach might be particularly sensitive to the period in which the Federal Reserve abandoned targeting the federal fund rate between 1979 and 1982. Therefore, in Section 5 we redo the analysis starting the sample in 1985, and showing that our results are not driven by these large monetary policy shocks in the early 80s.

[11] The responses for the difference between the 90th and the 10th percentile of the cross-sectional distribution and the IQR are reported in Figure 20. Given the very high correlation among dispersion measures, the IRFs display similar patterns differing mainly in the magnitude of the response.

Figure 4: Impulse responses of the year-on-year inflation rate as well as the median and the standard deviation of the individual inflation rate distribution

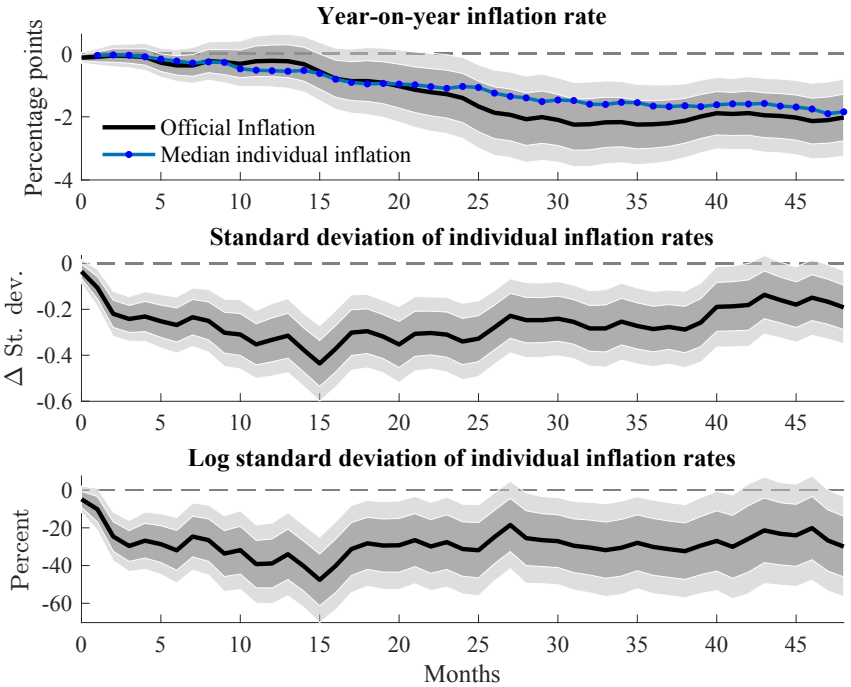

Notes: The figure plots in the top panel the impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the official annual inflation rate (black line) and the median inflation rate (blue line) of the individual inflation rate distribution. The middle panel reports the impulse response using as the dependent variable the dispersion in inflation, measured by the cross-sectional standard deviation and the bottom panel the log of the dispersion measure such that it can be interpreted as percent change relative to the steady state. The horizontal axis is in months. Impulse responses are computed at the monthly frequency using data for the period 1980M1:2007M12.

that monetary policy shocks lead to a decrease in the inflation dispersion in the economy.

To quantify the magnitude of the decrease in the inflation dispersion, the bottom panel computes the same impulse response but uses the log of the dispersion measure as the dependent variable, such that the magnitude can be interpreted as a percentage change relative to the steady state. Following a contractionary shock, we find that the cross-sectional standard deviation of inflation rates at the household level decreases

by around 40% after 3 years and approximately 20% at the end of the horizon considered. The average inflation rate over the time period considered is about 3.75% so a decrease of 1.5 percentage points corresponds to a decrease of 60% of the average value.

3.3 Importance of monetary policy shocks

Now, we assess the importance of monetary policy shocks for inflation dispersion. The forecast error variance decomposition, that is the share of the variance explained by the shock, is commonly used to answer this question.

For this purpose, we adopt the recently proposed estimator by Gorodnichenko and Lee (2019). Using as estimated forecast error $\hat{f}_{t+h,t-1}$ the residuals of the regression in equation (2), we then estimate the following equation:

$$\hat{f}_{t+h,t-1} = \alpha_0 e^{RR}_{t+h} + ... + \alpha_h e^{RR}_{t} + \tilde{v}_{t+h,t-1} \tag{3}$$

where e^{RR}_t is the shock at time *t* and $\tilde{v}_{t+h,t-1}$ is the error term due to innovations orthogonal to the shock series.

Our estimate of the share of the variance in dispersion explained by the shock is given by the R^2 of (3) which, by construction, is between 0 and 1. This measure provides an estimate of the extent to which monetary policy shocks are quantitatively important in driving dispersion dynamics.

The results from the variance decompositions are presented in Figure 5. Consistently with the impulse responses of the previous section, monetary policy shocks account for around 20% of the forecast error variance in the long run across the three measures of dispersion considered. These results are quantitatively in line with the contribution of monetary policy shocks to other inequality measures (Coibion et al.,

Figure 5: Forecast error variance decomposition for dispersion measures

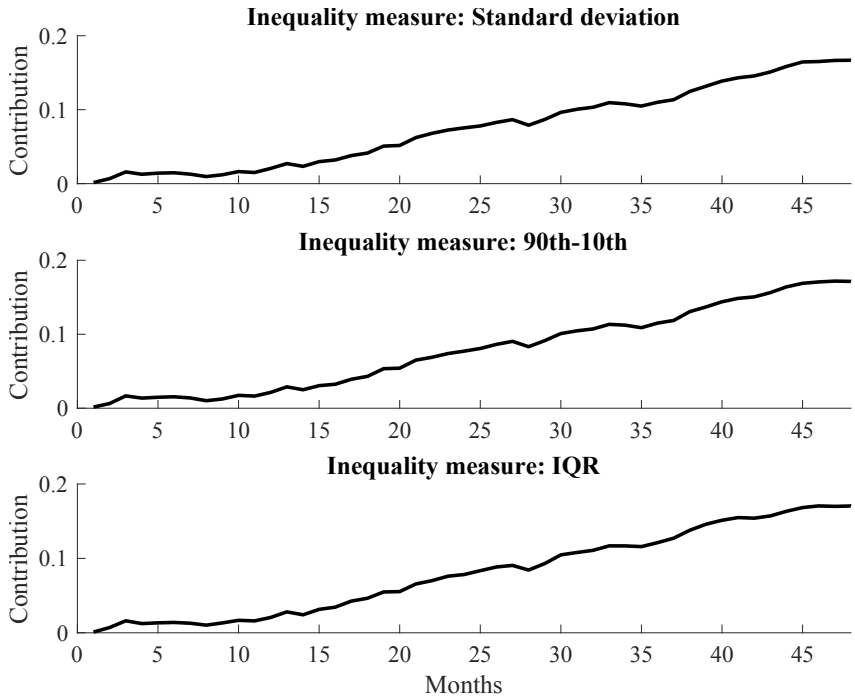

Notes: The figure plots the contribution of monetary policy shocks to the forecast error variance for the respective measure of inflation dispersion at different time horizons (in months).

2017, document that monetary policy shocks account for 10-20% of forecast error variance for expenditure and consumption inequality) as well as macroeconomic variables (Christiano et al., 1999).

3.4 Distributional consequences of monetary policy

In the previous section, we show how the dispersion of inflation decreases after a contractionary monetary policy shock. We now study more in-depth the main potential drivers of this result. Indeed, monetary policy can lead to a decrease in standard deviation through multiple channels.

We first evaluate how the two tails of the distribution react to monetary policy. Therefore, for every period we compute the change in inflation for the households at the 90th, 50th, and 10th percentile in terms of individual inflation rates. We then assess how the distance between the inflation at the 90th and the 50th, as well as the 50th and the 10th percentiles, respond to monetary policy shocks.

Figure 6: Distributional effects of monetary policy shock

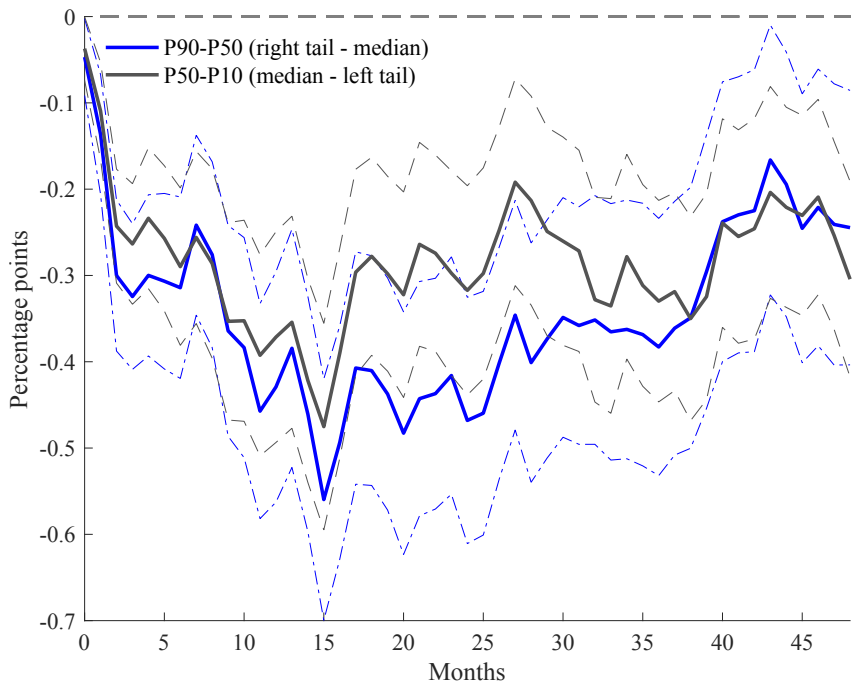

Notes: The figure plots impulse responses to a one percentage point contractionary monetary policy shock with the one standard deviation confidence intervals for the absolute distance between the inflation at the 95th and 50th percentile (blue line) as well as between the 50th and the 5th percentile (black line). The horizontal axis is in months. Impulse responses are computed at the monthly frequency using data for the period 1980M1:2007M12.

Results are reported in Figure 6. After a contractionary shock, the distance between both tails and the median decreases significantly. However, the responses are asymmetrical: at the horizon, the magnitude for the distance between the 90th and the 50th percentiles is more than twice the magnitude for the difference between

the 50th and the 10th percentiles. This implies that part of the decrease in inflation dispersion previously observed is driven by a stronger reduction in the inflation rate for the households at the top of the distribution.

Combining the results obtained, we can understand how monetary policy shocks shape the cross-sectional inflation distribution. Following a contractionary shock, the right tail of the distribution converges much more strongly towards the median, when compared to the left tail. Not surprisingly, since both distances decrease, we observe a reduction in the overall standard deviation. Finally, the tails behavior suggests as well that the mass of households on the right side of the median will squeeze much more than on the left side predicting that the distribution will be more negatively skewed.

In Figure 7 we test this hypothesis by estimating equation (2) using the overall skewness of individual inflation as the dependent variable. As expected, a contractionary monetary policy shock results in a significant and persistent decrease in the overall skewness of the distribution.

To ease the understanding of the previous results, Figure 8 illustrates the changes in the cross-sectional inflation distribution caused by a monetary policy shock. We start by plotting a Pearson distribution where the first three moments are equal to the unconditional mean, standard deviation, and skewness of the individual inflation rates distribution (black solid line). It approximates the empirical distribution. We then derive the impulse responses of the inflation rate for the different deciles level of the individual inflation distribution and retrieve the resulting mean, standard deviation, and skewness 48 months after the shock to compute the final distribution (black dashed line). The blue dashed line is obtained by simply shifting the black dashed line such that the median is kept constant to better visualize the changes in the second and third moments.

Figure 7: Impulse responses of inflation skewness

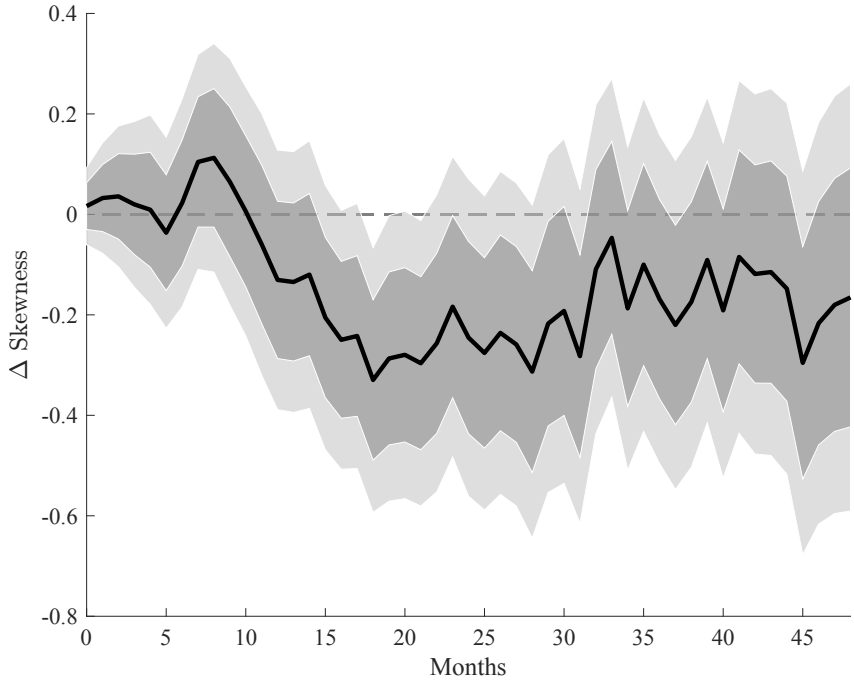

Notes: The figure plots impulse response to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the skewness of the inflation distribution. The horizontal axis is in months. Impulse responses are computed at the monthly frequency using data for the period 1980M1:2007M12.

A contractionary shock shifts the entire distribution to the left since it reduces aggregate and individual inflation rates. In line with the results from Figure 6, the distance between the right tail and the median reduces more than the distance between the left tail and the median. This has two consequences: first of all, the overall dispersion decreases. Secondly, the asymmetric response of the inflation rates at the two tails makes the distribution more right-skewed (i.e. the skewness of the distribution decreases).

Figure 8: Graphical representation of the cumulative change in individual inflation after a monetary policy shock

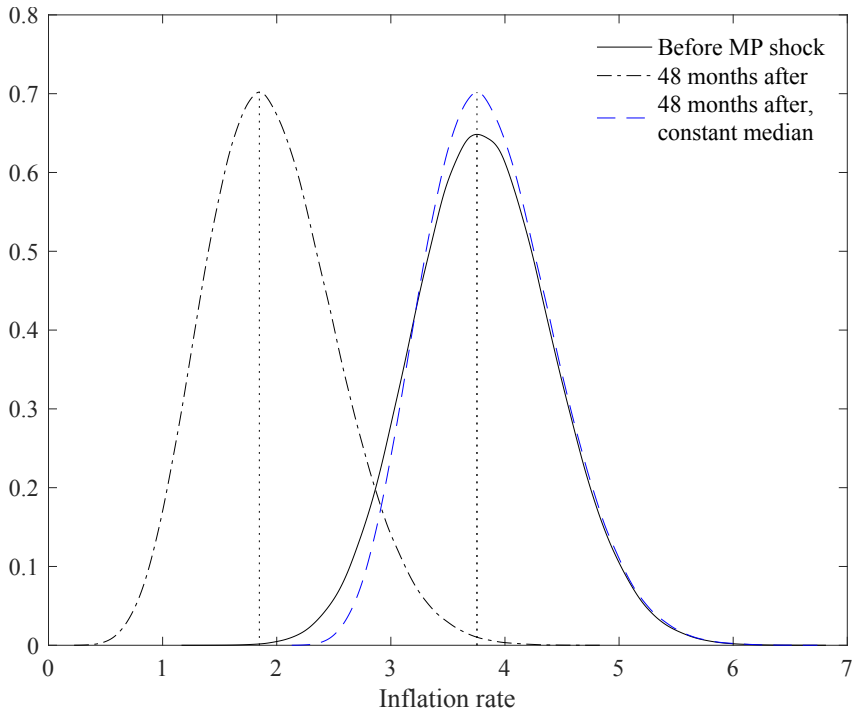

Notes: The figure shows a Pearson distribution of individual inflation rates before (black solid line) a contractionary monetary policy shock using as mean, standard deviation, and skewness the unconditional moments across time. The black dashed line reports the distribution of individual inflation rate 48 months after a contractionary monetary policy shock, matching the first three moments as calculated in the previous section. The blue dashed line displays the same distribution, keeping the median constant to better underline the changes in the second and third moments.

4 Heterogeneity across demographic groups

Having shown that monetary policy shocks decrease inflation dispersion in the economy, we now evaluate which observable household characteristics are more sensitive to monetary policy shocks in terms of individual inflation and how this affects the cross-sectional inflation dispersion. We focus in particular on three demographic groups: expenditure, salary, and income deciles.

4.1 Expenditure weights

Heterogeneity in inflation rates comes from the fact that households consume different consumption bundles. We show average expenditure weights for the first, fifth, and tenth decile of income, salary, and expenditure deciles for each of the 21 categories in table 2.

Several interesting facts can be noticed: first of all, the pattern across deciles is quite similar for income, salary, and expenditures. This already anticipates that the decile-level inflation rates of these three categories will react in a consistent way to monetary policy shocks. Second, although the weight for most of the categories either decreases or increases from the first to the tenth deciles, some categories display a U-shape pattern (e.g. *Gasoline*, *Medical expenses*). This is consistent with the findings of Cravino et al. (2020) who document that the highest price volatility is experienced by middle-income households. Finally, from the differences in weights across deciles, we can understand which categories are mainly responsible for inflation inequality. In particular, as we will show, a major role is played by *Food at home* and *Energy*. Indeed, not only low and middle-income households consume a significantly higher share of their income on these categories with respect to high-income households but their price level is also extremely volatile.

4.2 Impulse responses by demographic groups

To evaluate how different demographic groups react to monetary policy shocks, we start by estimating the LP with R&R shocks using as the dependent variable the cross-sectional standard deviation of the median inflation rates across expenditure, salary, and income deciles which we defined as *inflation inequality*[12].

[12] Appendix C explains in detail how the median inflation rates are computed following the same approach adopted by the BLS.

Figure 9: Impulse responses of inflation dispersion across expenditure, salary, and income deciles

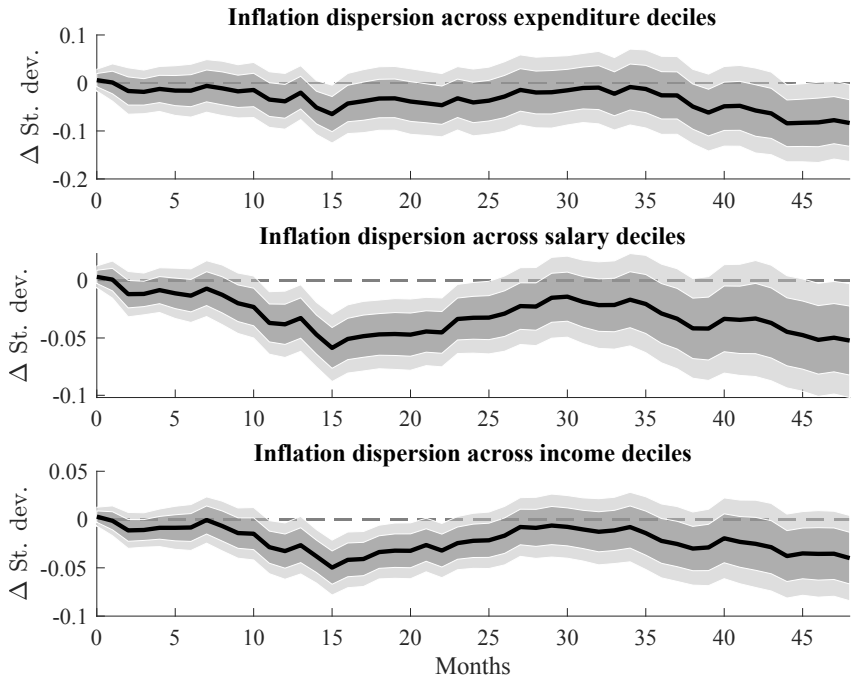

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for inflation dispersion across expenditure (top), salary (middle), and income deciles (bottom). The horizontal axis is in months. Dispersion is measured using the cross-sectional standard deviation. Impulse responses are computed at the monthly frequency using data for the period 1980M1:2007M12.

As one can see from Figure 9, following a contractionary monetary policy shock inflation inequality for the three categories significantly and persistently decreases. To better understand the main drivers of this result, we focus first on the median inflation rates for the different expenditure deciles, whose relative impulse responses are reported in the left panel of Figure 10.

Similar to what Cravino et al. (2020) found for income, the annual inflation rate of the households at the top of the expenditure distribution reacts substantially less to monetary policy shocks than the one in the middle. This relationship can be better visualized by looking at the right panel of Figure 10 which shows the response

Figure 10: Impulse responses across expenditure deciles

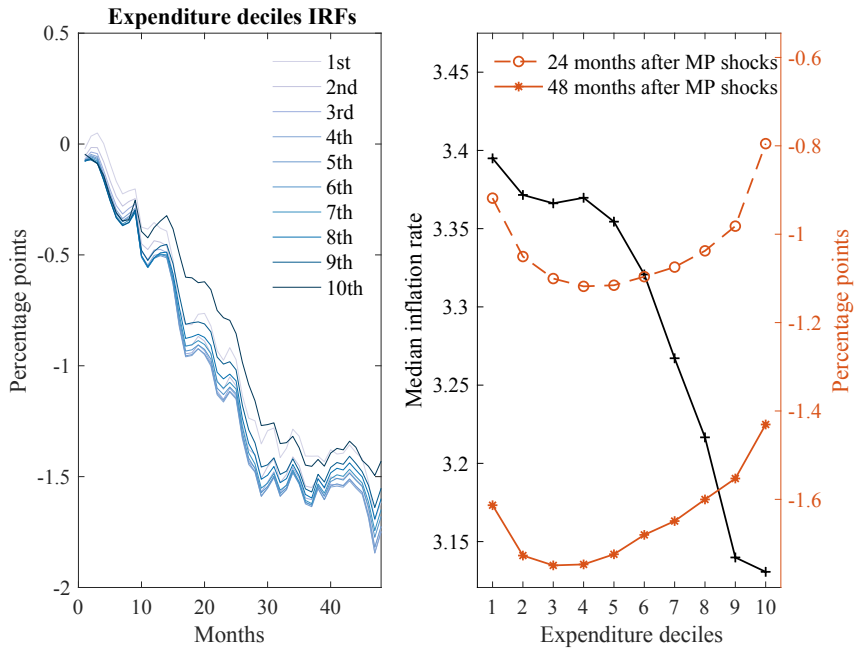

Notes: The left figure plots the impulse responses of the different expenditure deciles following a one percentage-point increase in the monetary policy shock, with monthly data for the period 1980M1:2007M12. The right plot reports the median inflation rate across deciles (left axis) as well as the cross-sectional distribution of the responses after 24 and 48 months after the shock (right axis).

across deciles after 24 and 48 months (dashed and dotted red lines, right axis). The difference between middle- and high-expenditure households is economically sizable. After 24 months, the annual inflation rate of the households in the top decile responds around 40% less than the inflation rate of the households in the fifth decile. After 48 months, the difference is still around 25%.

How does this relate to inflation inequality? We report in the same panel the median inflation rates across expenditure deciles relative to the time period considered (black line, left axis). One can notice how the higher the decile the lower the median inflation rate. While this result is not new per se[13], most of the literature

[13]See, among others, Jaravel (2019) and Kaplan and Schulhofer-Wohl (2017).

Figure 11: Impulse responses across salary deciles

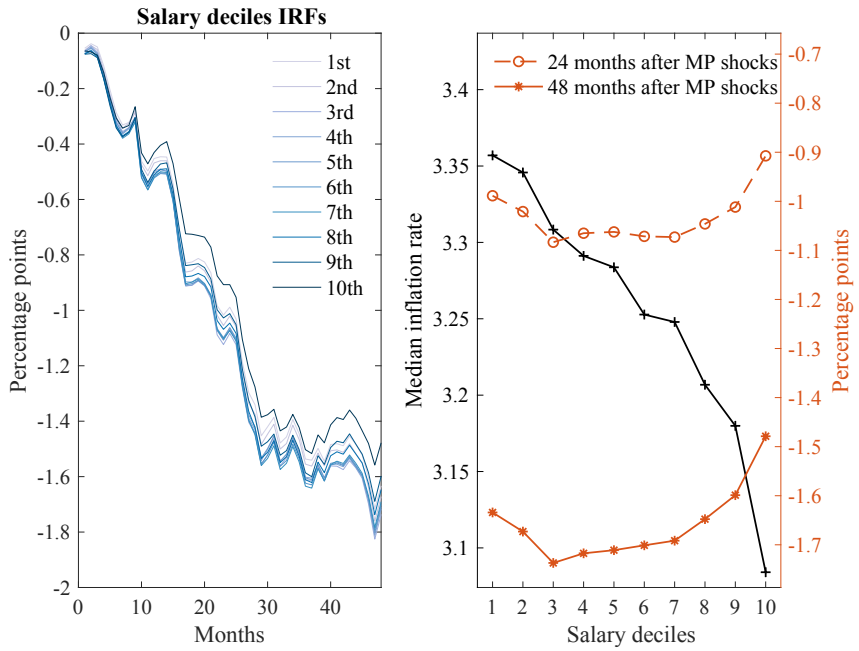

Notes: The left figure plots the impulse responses of the different salary deciles following a one percentage point increase in the monetary policy shock, with monthly data for the period 1980M1:2007M12. The right plot reports the median inflation rate across deciles (left axis) as well as the cross-sectional distribution of the responses after 24 and 48 months after the shock (right axis).

that focused on inflation inequality used the Nielsen scanner data available only from 2004. Here we show that the heterogeneity in inflation rates across income groups holds considering the period 1980-2007 as well. Therefore, on the one hand, given their consumption bundle, the high-expenditure households experience a lower median inflation rate than the households on the left side of the distribution. On the other hand, their inflation rate reacts significantly less to monetary policy shocks. These two results combined imply that following a contractionary shock, we observe a convergence of individual inflation rates across the distribution leading to a lower inflation inequality as documented in Figure 9. Similar results can be found focusing on salary and income deciles as shown in Figure 11 and Figure 12 respectively.

Figure 12: Impulse responses across income deciles

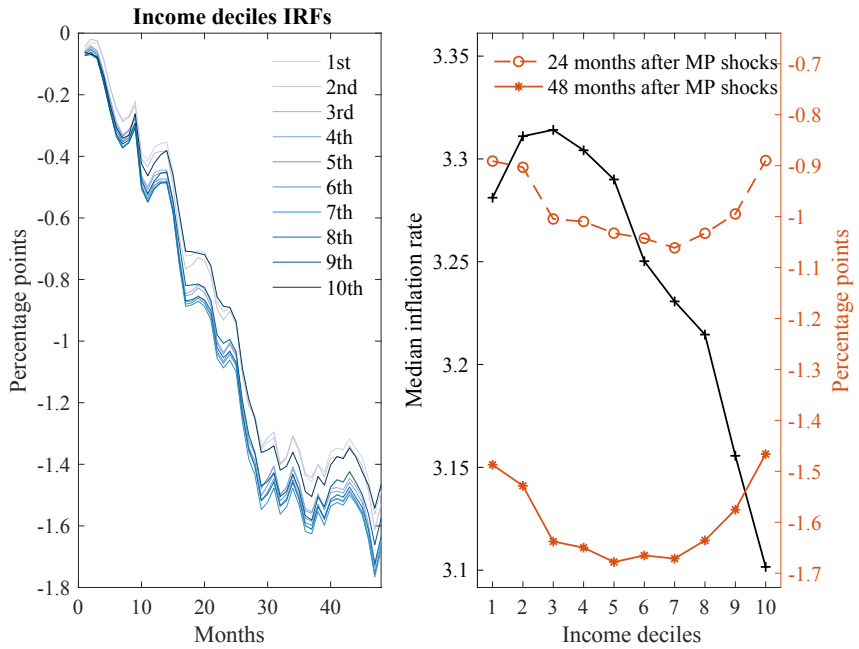

Notes: The left figure plots the impulse responses of the different income deciles following a one percentage point increase in the monetary policy shock, with monthly data for the period 1980M1:2007M12. The right plot reports the median inflation rate across deciles (left axis) as well as the cross-sectional distribution of the responses after 24 and 48 months after the shock (right axis).

To summarize, our empirical analysis strongly suggests that monetary policy shocks can have significant and non-negligible distributional effects on the economy. Since the inflation rate of higher-income households is lower in median relative to lower- and middle-income deciles, and at the same time their inflation rate is less reactive to unexpected changes in the interest rate, this leads to a decrease in inflation inequality following a contractionary shock.

4.3 Real expenditure inequality

Does the identified inflation inequality have any effect on the estimated impact of monetary shocks on real expenditure inequality? To answer this question, we follow

Coibion et al. (2017) as close as possible and compute a broad measure of household expenditure which includes non-durables, durables, and services[14]. Few expenses are excluded since the relative sub-category price index is not easily identifiable (e.g. occupational expenses, mortgage, and property taxes).

To evaluate the role played by inflation inequality, we create two different series for real expenditure. On the one hand, in line with the literature, we deflate each category by the aggregate CPI-U. On the other hand, we deflate each item group by its relative price index. We then aggregate the expenditures at quarterly levels to reduce sampling error and to avoid that unusual purchases might bias the analysis. The measure of expenditure is winsorized at the bottom and top 1 percent. Inequality across households is computed as the cross-sectional standard deviations of log levels, the Gini coefficients of levels, and the difference between the 90th percentile and the 10th percentile of log levels. Finally, all series are seasonally adjusted.

Expenditure inequality is defined as $Ineq_t^{IH}$ and $Ineq_t^{NoIH}$ respectively for when inflation heterogeneity is taken into account by deflating each category by the relative price index and for when it is neglected. As an example, the standard deviations at time t across households i is equal to $Std(logC_{i,t}^{IH})$ and $Std(logC_{i,t}^{NoIH})$ with:

$$C_{i,t}^{IH} = \sum_{j \in J} \frac{C_{i,j,t}}{P_{j,t}}, \qquad C_{i,t}^{NoIH} = \sum_{j \in J} \frac{C_{i,j,t}}{P_t} \tag{4}$$

where $C_{i,j,t}$ is the nominal consumption of household i relative to category j at time t, $P_{j,t}$ is the price index of the category j at time t and P_t is the aggregate price index.

[14]In particular, the categories considered are: Food at Home, Food Away, Alcohol at Home, Alcohol Away, Apparel, Gasoline, Personal Care (services and durables), Reading, Tobacco, Household Furnishings and Operations, Energy, Water, Other Lodging, Public Transportation, House expenditures (services and durables), Rental expenditures (services and durables), Rent paid, Heath insurance, Health expenditures (services and durables), Education, Vehicles purchase, Vehicle expenditures (services and durables), Miscellaneous.

Figure 13: Impulse responses of expenditure inequality

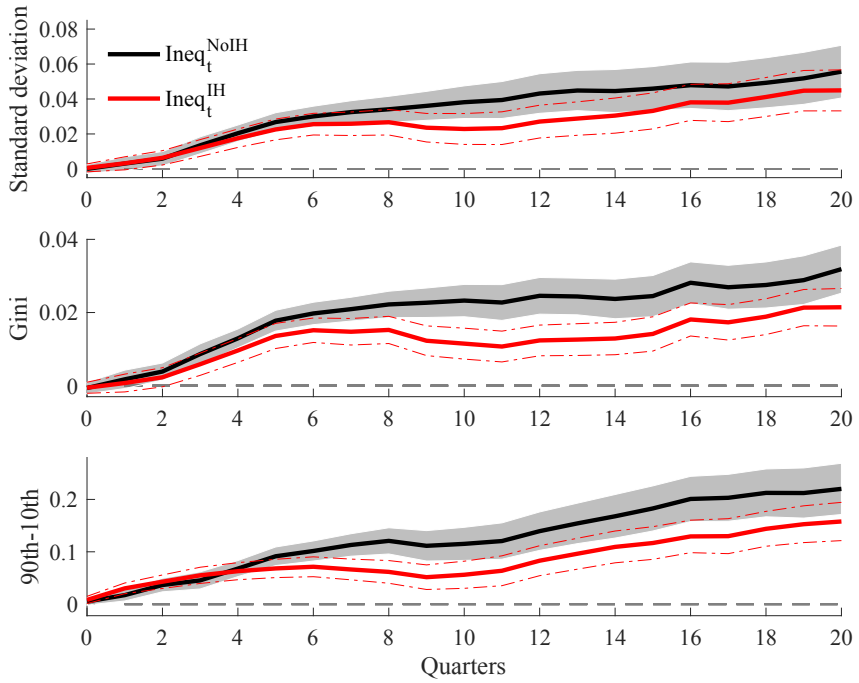

Notes: The figure plots impulse responses to a one percentage point contractionary monetary policy shock, as well as one standard deviation confidence intervals for expenditures inequality. The horizontal axis is in quarter and inequality is measured using the cross-sectional standard deviation (left), Gini coefficient (middle), and the log difference between the 90th and 10th percentiles of the cross-sectional distribution (right). The black solid line and the dark grey shaded areas depict the impulse response obtained deflating the expenditure categories by the aggregate CPI, the red solid line and the dashed red lines refer to the impulse obtained by deflating each category by their respective price index. Impulse responses are computed at the quarterly frequency using data for the period 1980Q1:2008Q4.

In order to make our results as comparable as possible, we use the same econometric procedure adopted by Coibion et al. (2017) (i.e. local projection with Romer and Romer, 2004 shocks at quarterly frequency) over the same time period, 1980Q1:2008Q4[15]. Since the series are quarterly, we include as controls 20 lags for the shocks and 2 lags for the dependent variable and we compute the impulse responses over 20 quarters.

[15]Similar results are obtained restricting the analysis to our baseline period.

The results are reported in Figure 13. The black solid lines report the impulse responses of the three measures of expenditure inequality obtained deflating the expenditure categories by the aggregate CPI. The shape and the magnitude of the responses are very close to the ones obtained by Coibion et al. (2017). After a contractionary monetary policy shock, expenditure inequality persistently and significantly increases. However, neglecting inflation heterogeneity across consumption bundles leads to an overestimation of the overall effect. As shown by the red solid lines which report the responses of the expenditure inequality measures obtained by deflating each category by their respective price index, when the expenditure categories are properly deflated, the estimated effect of monetary policy on inequality is approximately 20% lower for standard deviation and 30% for the Gini coefficient and the 90th-10th percentile difference. It is worth mentioning that the estimated coefficients are still positive and significant, such that monetary policy still has redistributive effects on the economy.

This result can be explained by combining the new empirical evidence of the previous sections. Along the income distribution, a contractionary monetary shock has heterogeneous effects on nominal consumption. The nominal consumption of low and middle-income households decreases more than that of high-income households because they are more sensitive to the monetary policy shock (e.g. they are more likely to lose their job in an economic downturn). However, at the same time, the cost of their consumption basket decreases more strongly as well. Hence, the effect is partially offset in real terms. Real consumption heterogeneity, therefore, increases less than when a common inflation rate is assumed, resulting in a more muted, but still significant, response of expenditure inequality.

5 Robustness

In order to strengthen the validity of our findings in the previous sections, we show that our results are robust across a wide range of alternative specifications. Firstly, we compare the results using the R&R shocks as instrument for the change in interest rate (IV-LP) instead of directly inserting them into the LP. Secondly, we assess the sensitivity of our results to different lag specifications. Thirdly, we perform the same analysis starting our sample in 1985M1 to control for the Volcker disinflation period. Finally, we show that the number of categories we chose to compute inflation rates at household levels played no role in our results. The figures are reported in Appendix E.

5.1 Local projection with and without instrumental variables

The LP approach assumes that the shock series captures the true monetary shock. However, a more reasonable assumption may be to merely assume that the R&R shocks are correlated with the true monetary shocks and uncorrelated with other structural shocks. For a full discussion of the difference, see Ramey (2016) and Stock and Watson (2018).

We can then recover the impulse response to a monetary shock from LP and two-stages least squares where the US federal fund rate is instrumented with the R&R shocks series (IV-LP). Specifically, in the first stage, we regress the R&R shocks on the change of the interest rate and in the second stage, we estimate the LP using, instead of the shock series itself, the fitted value of the first stage regression, as discussed in Stock and Watson (2018):

$$\Delta i_{t,t-1} = c^1 + \beta^1 e_t^{RR} + \gamma^1(L) controls_t + \mu_t$$

$$x_{t+h} - x_{t+h-1} = c_h^2 + \beta_h^2 \Delta \hat{i_{t,t-1}} + \gamma_h^2(L) controls_t + \epsilon_{t+h}$$

for $h = 1, ..., 48$. As in the main analysis, e_t^{RR} is the R&R shocks series, $\Delta i_{t,t-1}$ is the change in federal fund rate and we control for 6 lags of the dependent variable and 48 lags of the change in interest rate.

As a further robustness check, we replicate the main results of our empirical analysis using the IV-LP approach with R&R shocks. The results for the median and cross-sectional standard deviation of the individual inflation rates are reported in Figure 22 and Figure 23 shows the responses using as the dependent variable the difference between the 90th and the 10th percentile of the cross-sectional distribution and the IQR. The magnitude of the impulse responses is slightly higher when compared to the LP approach but the shape of the impulse response is unaffected. With this specification as well, we can statistically claim that contractionary monetary policy shocks have redistributive effects in terms of inflation dispersion.

The results from the variance decompositions using IV-LP (Figure 24) confirm as well that monetary policy shocks account for around 20% of the forecast error variance in the long run across all the three measures of dispersion considered.

Figure 25 shows how the different deciles for income, salary, and expenditure are affected by monetary policy shocks. The size of the coefficients is again higher than with respect to the baseline specification, but the main results still hold: contractionary monetary policy shocks significantly reduce inflation inequality.

Finally, Figure 26, Figure 27, and Figure 28 report the impulse responses for median inflation rates of the different deciles as well as the unconditional median inflation rate. Even in this case, similarly to what we found with the LP approach, households in higher deciles experienced a lower inflation rate across time. Moreover, their median inflation is less sensitive to monetary policy shocks therefore it

decreases less than the one of lower deciles. As a consequence, inflation inequality persistently decreases after a contractionary monetary policy shock.

5.2 Different lag specification

We re-estimate equation (2) again with alternative lag specification. In Figure 21 we run the LP regression including all possible combinations between 24 and 60 lags for the monetary policy shocks as well as 5 and 12 lags for the cross-sectional standard deviation of the individual inflation alongside our baseline specification. Similar results are obtained for the other measures of dispersion as well. Changing the number of lags has insignificant effects on the impulse responses: after a contractionary monetary policy shock, inflation dispersion significantly decreases.

5.3 Volcker disinflation

Coibion (2012) showed how few episodes in the early 80s can be the main drivers of the impulse responses computed using LP with R&R shocks. Since then, it has been common practice for researchers to test their results excluding the period between 1979 and 1982 in which the Federal Reserve abandoned targeting the federal fund rate. Figure 29 reports the IRFs obtained using the baseline specification but starting the sample in 1985M1. In this case, the results are robust as well.

As one can see in Figure 30, the results for the difference in inflation rates across deciles are largely unaffected even excluding the Volcker disinflation period by starting the sample in 1985M1.

5.4 Alternative aggregation

Finally, we test whether the expenditure categories we chose played any role in our results. The BLS reports the CPI at different granular levels. Not only do we have data for *Food and Beverage*, the most aggregate item category, but we also for the sub-category *Eggs*, the most disaggregate. In choosing the baseline aggregation, we faced a trade-off between using data as disaggregate as possible to fully capture inflation dispersion and the quality of the price index. Not all price series are indeed available since the early 80s and this is true especially for the most disaggregate goods and services.

Following the literature, we adopt a rather conservative aggregation. In this section, we show that using alternative expenditure categories, the main results are basically unaffected.

In Figure 31, we report the response from our baseline specification with 21 categories (blue line) against three alternative aggregations: using price indices at a much more granular level[16] (top left), an even more conservative number of categories[17] (top right), and an intermediate aggregation[18] (bottom left).

[16] 31 categories: Food at Home, Food Away from Home, Alcohol, Rental expenditures (durables), Rental expenditures (services), Rent Paid, Rent Equivalent, House Expenditures (durables), House Expenditures (services), Other House related expenses, Other Lodging, Energy, Water, Phone, Household Furnishings and Operations, Jewelry, Clothing (durables), Clothing (services), Gasoline, Vehicles Expenditures (durables), Vehicles Expenditures (services), Public Transportation, Medical, Entertainment, Personal Care (durables), Personal Care (services), Reading, Education, Tobacco and Other Expenses.

[17] 14 categories: Food, Alcohol, Housing, Apparel, Gasoline, Other Vehicle Expenses, Public Transportation, Medical, Entertainment, Personal Care, Reading, Education, Tobacco and Other Expenses.

[18] 27 categories: Food at Home, Food Away from Home, Alcohol at Home, Alcohol Away from Home, Owned Dwellings, Rented Dwellings, Other Lodging, Energy, Water, Phone, Household Furnishings and Operations, Jewerly, Clothing (durables), Clothing (services), Gasoline, Vehicles Expenditures (durables), Vehicles Expenditures (services), Public Transportation, Medical, Entertainment, Personal Care (durables), Personal Care (services), Reading, Education, Tobacco and Other Expenses.

As one can notice, the magnitude of the responses is rather similar to the one obtained in our baseline specification. Increasing the number of categories slightly increases the size of the response but the change is negligible. Reducing the number of categories does not affect our estimates.

This result is particularly important because it partially allows us to address one of the potential critic of the paper: the substitution bias across expenditure categories. Throughout the paper, we have conducted our analysis under the assumption that changes in inflation dispersion are driven mainly by changes in prices and that changes in expenditure shares play only a marginal role. Cravino et al. (2020) already tested this assumption for CEX by using the differences between the Laspeyres and Paasche price index as a proxy for the substitution bias from 1987 and 2004. The authors showed that the difference between the two indices is negligible over time demonstrating that the substitution bias must be very small.

Furthermore, using the Nielsen data, Jaravel (2019) evaluates whether the observed inflation heterogeneity along the income distribution stems from the fact that high-income households purchase different goods or whether they pay more for the same goods, for instance, because they buy from different shops. The inflation difference is then decomposed into a *between* and a *within*-component. The former corresponds to the inflation difference that we would observe if households differ only in terms of the expenditure shares across categories and experience the same within-category inflation. Vice versa, the latter refers to the difference that would arise in case the households experience the same within-category inflation but had different expenditure shares. The between component accounts for more than 70% of the inflation difference.

The three plots in Figure 31 seem to confirm these findings. The size of the response for inflation dispersion following a contractionary monetary policy shock is basically unaffected when we increase and decrease the number of expenditure

categories. This suggests that inflation dispersion is mainly caused by heterogeneity in prices across rather than within expenditure categories.

Finally, the bottom right plot shows the inflation dispersion response when we exclude the categories *Food*, *Energy* and *Gasoline*. This new specification is close to the definition of Core CPI that the Federal Reserve Bank uses to decide which monetary policy to adopt. Not surprisingly, removing three of the most volatile categories cancels out the response of inflation dispersion almost entirely.

6 Conclusion

Macroeconomic research has shifted focus towards the heterogeneity of consumers and inequality. We provide a refined view of the relationship between the two by exploiting the differences in individual inflation rates, which households face due to different consumption bundles.

To do so, we use individual expenditure data from the CEX and combine it with category-level inflation rates from the CPI to get individual inflation rates and dispersion measures. We then evaluate monetary policy transmission by studying the response of these measures to Romer and Romer (2004) shocks. We show that contractionary monetary policy significantly and persistently leads to lower levels of inflation dispersion across households. On a five-year horizon, monetary policy accounts for approximately 20% of the variation in dispersion. This effect is not symmetric. Indeed, the right tail of the distribution converges more strongly to the median than the left tail, leading to a slightly left-skewed distribution.

We show that the median inflation rate of the high-income households is significantly less reactive to contractionary shocks than the one of the lower- and middle-income households, leading to an overall reduction of inflation inequality across income groups. The same is true for expenditure and salary deciles.

These findings contribute to a refined, more accurate view of consumption and income differences across households because they account for differences in inflation rates faced by consumers. They are relevant because the differences between households are sizable and non-random. They are systematically different across demographic groups, and they react differently to monetary policy shocks. Hence, assuming that all consumers experience the same inflation rate in a heterogeneous agent setting may lead to biased results, as well as exaggerated levels of inequality. With a more accurate picture of the purchasing power of individual consumers, future researchers will be able to more accurately measure the real economic conditions faced by individual households.

References

Barakchian, S. M. and Crowe, C. (2013). "Monetary policy matters: Evidence from new shocks data". *Journal of Monetary Economics, Vol. 60, pp. 950–966.*

Bee, A., Meyer, B., and Sullivan, J. (2013). "The Validity of Consumption Data: Are the Consumer Expenditure Interview and Diary Surveys Informative?". In *Improving the Measurement of Consumer Expenditures*, NBER Chapters, pages 204–240. National Bureau of Economic Research, Inc.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (1999). "What have we learned and to what end?". *In J. B. Taylor and M. Woodford (Eds.), Handbook of Macroeconomics, Volume 1A, Chapter 2, pp. 65-148. North-Holland*.

Cloyne, J., Ferreira, C., and Surico, P. (2019). "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism". *The Review of Economic Studies, Volume 87, Issue 1, Pages 102–129*.

Coibion, O. (2012). "Are the Effects of Monetary Policy Shocks Big or Small?". *American Economic Journal: Macroeconomics 2012, 4(2): 1–32*.

Coibion, O., Gorodnichenko, Y., Kueng, L., and Silvia, J. (2017). "Innocent Bystanders? Monetary policy and inequality". *Journal of Monetary Economics, 2017, vol. 88, issue C, 70-89*.

Cravino, J., Lan, T., and Levchenko, A. (2020). "Price Stickiness Along the Income Distribution and the Effects of Monetary Policy". *Journal of Monetary Economics, 110:19-32.*

Driscoll, J. and Kraay, A. (1998). "Consistent covariance matrix estimation with spatially dependent panel data". *Review Econ. Stat. 80, 549–560*.

Fisher, D. J., Johnson, S. D., and Smeeding, M. T. (2013). "Measuring the Trends in Inequality of Individuals and Families: Income and Consumption". *American Economic Review 103(3), 184-88*.

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

Gorodnichenko, Y. and Lee, B. (2019). "Forecast Error Variance Decompositions with Local Projections". *Journal of Business and Economic Statistics*.

Hobijn, B. and Lagakos, D. (2005). "Inflation inequality in the United States". *Review of Income and Wealth*.

Jaravel, X. (2019). "The unequal gains from product innovations: evidence from the U.S. retail sector". *The Quarterly Journal of Economics, Volume 134, Issue 2, Pages 715-783*.

Johannsen, B. K. (2014). "Inflation Experience and Inflation Expectations: Dispersion and Disagreement Within Demographic Groups". *FEDS Working Paper No. 2014-89*.

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

Kaplan, G., Moll, B., and Violante, G. (2018). "Monetary Policy According to HANK". *American Economic Review*, 108(3):697–743.

Kaplan, G. and Schulhofer-Wohl, S. (2017). "Inflation at the household level". *FRB of Chicago Working Paper No. WP-2017-13*.

Mumtaza, H. and Theophilopoulou, A. (2017). "The impact of monetary policy on inequality in the UK. An empirical analysis". *European Economic Review, Volume 98, Pages 410-423*.

Ramey, V. (2016). "Macroeconomic shocks and their propagation". *In Handbook of Macroeconomics. Vol. 2, 71*.

Romer, C. D. and Romer, D. H. (2004). "A new measure of monetary shocks: Derivation and implications". *American Economic Review 94(4), 1055-84*.

Samarina, A. and Nguyen, A. (2019). "Does monetary policy affect income inequality in the euro area?". *DNB Working Papers 626, Netherlands Central Bank, Research Department.*

Stock, J. H. and Watson, M. W. (2018). "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments". *The Economic Journal, Volume 128 Issue 610.*

Wong, A. (2019). "Refinancing and the Transmission of Monetary Policy to Consumption". *R&R American Economic Review, May 2019*.

For Online Publication

A Data sources

In this section, we document in greater detail the data sources used and the properties of the underlying data.

A.1 Price Indices

Since individual inflation rates are a weighted average of sectoral price indices, Table 1 displays the CPI subindices used, as well as their respective statistical properties.

Table 1: Item-level CPI statistics

CPI series (Item Code)[19]	Mean	Median	Standard Deviation	p10	p90
Food at Home (SAF11)	3.05	2.72	1.84	1.01	5.60
Food Away from Home (SEFV)	3.36	3.05	1.41	1.99	4.61
Alcoholic Beverages (SAF116)	3.24	2.73	1.84	1.64	5.29
Rented Dwellings (SEHA)	3.94	3.60	1.53	2.46	6.15
Owned Dwellings (SEHC)	3.65	3.33	1.01	2.42	5.13
Other Lodging (MUUR0000SE2102-SEHB)	5.15	4.65	3.50	1.51	9.69
Energy (SAH21)	3.29	2.41	5.74	-3.19	10.82
Water (SEHG01)	5.34	5.23	2.38	2.82	7.79
Phone (SAE2)	-1.06	-1.08	1.70	-3.31	1.13

[20]The official series ID, as defined by the BLS, is a combination of "CUUR0000", which stands for the unadjusted CPI-U inflation rate for the whole US, and the Item Code, as shown in the table.

Table 1: Item-level CPI statistics (continued)

CPI series (Item Code)[20]	Mean	Median	Standard Deviation	p10	p90
Household F&O[21] (SAH3)	1.43	1.34	1.77	-0.39	2.70
Apparel (SAA)	1.00	0.82	2.32	-1.83	4.49
Gasoline (SETB)	3.31	2.93	13.79	-13.63	20.98
Other Vehicle Expenses (SETB-SETD-SETE-SETF)	3.02	2.34	2.10	0.79	6.75
Public Transportation (SETG)	4.47	4.06	5.08	-0.93	9.54
Medical care (SAM)	5.72	4.82	2.21	3.45	9.01
Entertainment (SAR)	1.47	1.34	0.74	0.59	2.64
Personal Care (SAG1)	3.23	2.79	1.57	1.87	5.01
Reading (SERG)	3.64	3.36	2.50	0.86	7.01
Education (SAE)	2.40	2.40	0.96	1.10	3.70
Tobacco (SEGA)	7.56	7.11	6.08	2.27	12.75
Other Expenses (SEGD)	5.73	4.93	2.84	3.29	11.48
CPI-U (SA0)	**3.42**	**3.04**	**1.72**	**1.68**	**5.01**

A.2 Consumer expenditure survey data

In this section, we provide further details about the construction of the dataset we use in the empirical analysis. We download the raw data for the period 1980-2005 from the ASCII files available from the Inter-university Consortium for Political and Social Research (ICPSR) whereas from the year 2006 onward we use the data provided by the BLS. For each quarter, the Interview Survey is structured as follows: the expenditure data are recovered from the disaggregated MTAB files, income data

[21] Household Furnishings and Operations

are derived from the FMLY files and additional information regarding the households can be found in the MEMB files.

In line with the literature, we aggregate together expenditures about the same month which is reported in different interviews. Then, we drop households that report zero expenditure on food as well as those who report negative expenditures for categories that cannot be negative according to the data codebook, such as expenditures for elderly care. Respondents younger than 25 years and older than 75 are excluded as well. In order to correct for sample breaks caused by slight changes in the questionnaire (food at home (1982Q1-88Q1), food away from home (2007Q2), and personal care services (2001Q2)) we regress each expenditure series on a time trend and indicators for the corresponding sample breaks and then subtract from the original series the effect of the dummies. For all these transformations, we rely heavily on Coibion et al. (2017).

Finally, the CEX data started to include the imputed income since 2004. In order to impute income data before that year, we follow the approach adopted by Fisher et al. (2013) and Coibion et al. (2017): for households recording a bracket range, we use the median point of the bracket. Furthermore, we estimate the remaining income observations by regressing at the annual level and with sampling weights income on a set of observable characteristics such as age, age squared, the reference person's gender, race, education, number of weeks worked full or part-time in the last 12 month, unadjusted family size, the number of children less than 18, the number of persons over 64, the number of earners and using fixed effects for the income reporting date. To account for the sampling uncertainty, we add residuals drawn randomly with replacement from the sampling distribution to the predicted values. We then trim values above the top-coding threshold at the top coding value.

Expenditure weights

In a next step, we calculate expenditure shares from the cleaned expenditure data, which constitute the weights used to calculate individual inflation rates. The differences in these weights are highly consequential, as they are responsible for the dispersion in our dataset. We find substantial variation in the weights, variation that can be explained to a large part by either, income, salary, or expenditure deciles. For expositional purposes, Table 2 only shows the weights for the 1st, 5th, and 10th deciles.

Table 2: Expenditure weights for the first, fifth and tenth decile of income, salary, and expenditure

	Income deciles			Salary deciles			Expenditure deciles		
	1st	5th	10th	1st	5th	10th	1st	5th	10th
Food at Home	18.7	14.2	11.1	16.5	14.0	11.1	22.0	14.3	9.9
Food Away	7.2	7.5	7.3	7.7	7.6	7.2	8.0	7.3	6.9
Alcohol	1.0	1.1	1.2	1.1	1.2	1.2	1.1	1.1	1.1
Rented Dwellings	15.6	12.4	6.0	13.7	12.4	6.0	21.8	10.6	5.9
Owned Dwellings	15.4	17.1	22.6	14.5	16.8	22.8	6.5	19.3	22.6
Other Lodging	0.5	0.6	1.4	0.7	0.6	1.3	0.3	0.6	1.5
Energy	6.2	5.4	4.3	5.7	5.2	4.3	6.6	5.6	3.7
Water	0.9	1.0	0.9	0.9	0.9	0.9	0.9	1.0	0.8
Phone	3.4	3.0	2.3	3.2	3.0	2.3	3.8	3.0	2.1
Household F&O[22]	3.3	4.5	7.0	3.9	4.7	7.0	2.5	4.4	8.1
Apparel	4.0	4.3	5.6	4.4	4.6	5.7	3.7	4.2	5.7
Gasoline	4.2	5.3	4.4	5.0	5.6	4.5	4.3	5.4	3.8
Other Vehicle Expenses	4.3	6.8	7.2	5.5	7.2	7.3	3.7	6.8	7.0
Public Transportation	1.0	1.0	1.7	1.1	1.0	1.6	1.0	0.9	1.8
Medical	5.0	6.2	5.0	5.4	5.2	4.6	4.9	6.2	5.6

[22]Household Furnishings and Operations

Table 2: Expenditure weights for the first, fifth and tenth decile of income, salary, and expenditure (continued)

	Income deciles			Salary deciles			Expenditure deciles		
	1st	5th	10th	1st	5th	10th	1st	5th	10th
Entertainment	3.8	4.9	6.5	4.5	5.2	6.5	3.4	4.8	7.0
Personal Care	0.9	1.0	1.0	1.0	1.0	1.0	0.9	1.0	0.9
Reading	0.4	0.5	0.6	0.5	0.5	0.6	0.4	0.5	0.6
Education	1.6	0.8	2.3	2.1	1.0	2.4	1.3	0.9	2.9
Tobacco	1.7	1.4	0.6	1.7	1.4	0.6	2.3	1.3	0.6
Other Expenses	0.8	1.1	1.1	0.9	1.1	0.9	0.6	1.0	1.5

A.3 Matching of expenditure and inflation data

We matched the expenditure categories with the respective price indices. Following Hobijn and Lagakos (2005), for the category *Other Vehicle Expenses* which did not have a perfect match with the available CPI sub-categories, we create the CPI index by combining the series that match this category (that is, SETB, SETD, SETE, and SETF). For each period we used the official weights provided by the BLS, as displayed in the table "Relative Importance in the CPI". Finally, since *Other Lodging* changed the name, we used *Lodging away from home* until 1997 (MUUR0000SE2102) and *Lodging while out of town* (SEHB) until the end of the sample. In all cases, the CPI series we use are the not-seasonally-adjusted *US City Average for all urban consumers* series.

Table 3: Matching between CEX expenditure category and CPI

CEX Expenditure Category	CPI Series (Item Code)
Food at Home	SAF11
Food Away from Home	SEFV
Alcohol	SAF116
Owned Dwellings	SEHC
Rented Dwellings	SEHA
Other Lodging	MUUR0000SE2102-SEHB
Energy	SAH21
Water	SEHG01
Phone	SAE2
Household Furnishings and Operations	SAH3
Apparel	SAA
Gasoline	SETB
Other Vehicle Expenses	SETB-SETD-SETE-SETF
Public Transportation	SETG
Medical	SAM
Entertainment	SAR
Personal Care	SAG1
Reading	SERG
Education	SAE
Tobacco	SEGA
Other Expenses	SEGD

B Energy sector

Particularly interesting in Figure 3 are the sharp increases in dispersion in 1986 and the high volatility from 2000 onward. One possible explanation regarding these two episodes can be found in Figure 14. The official CPI and the inflation rate for the sub-category *Household Energy*, which includes expenditures like natural gas, electricity, and fuel, are depicted on the left axis. On the right axis, we report the inflation dispersion measured as the cross-sectional standard deviation.

Figure 14: Historical series of CPI, *Household Energy* inflation and the standard deviation of inflation dispersion

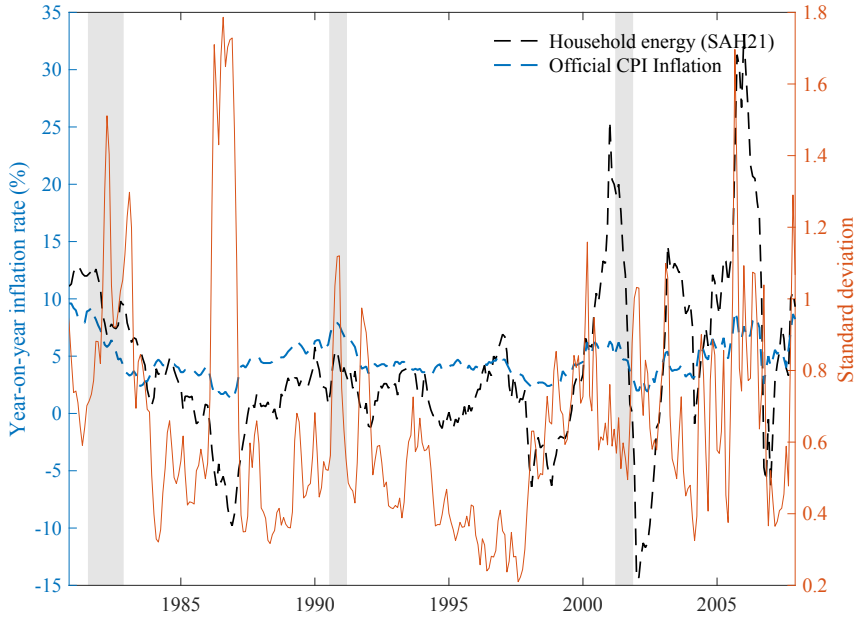

Notes: The official CPI (blue line) and the sub-category *Household Energy* (dashed black) are depicted on the left axis, whereas the evolution of inflation dispersion (dotted red), as measured with the cross-sectional standard deviation, is on the right axis. The gray shaded areas depict U.S. recessions.

The sub-category *Household Energy* has a weight of approximately 6% in the computation of the CPI, so its influence on aggregate rates is moderate. However,

at the individual household level the expenditure share is highly heterogeneous, especially along the income distribution (more on that in Subsection 4). As one can notice, the sharp drop in energy prices in 1986 and the higher volatility from 2000 onward partially explain the observed behavior in inflation dispersion.

C Decile-level expenditure weights

Households are interviewed a different number of times and for at most four consecutive quarters which corresponds to twelve months worth of spending information. However, this does not necessarily match with the calendar year. To control for this, we compute decile-based inflation rate closely following the BLS procedure as in Cravino et al. (2020). First of all, we sort households into deciles based on their annual income, salary, median, and mean expenditure. We then compute the average expenditure for each item category at every decile in the calendar year. For instance, a respondent interviewed in February will report its consumption for January but also for November and December of the previous year. Similar to what the BLS does for the computation of the official CPI, to account for the relative contribution of each household to the decile-mean value of a calendar year, we weight the consumption by the number of months a household reports expenditures during a calendar year (the BLS calls this variable MO_SCOPE).

We can then use the formula below to compute the average expenditure for each category j at each decile d. First, for household i at decile d, we aggregate over all the expenditures on good j during the calendar year. Second, the household total expenditure is weighted by the sampling weights, fwt, provided by BLS to make the survey sample representative of the U.S. population. Then, the weighted household expenditures are summed up at the decile-level. Finally, to obtain the monthly average income spent on good j by decile d, we divide the annual weighted household expenditures for category j by the weighted number of months household at decile d reported expenditures during the calendar year. To annualize the average category expenditure at the decile level, it is enough to multiply the monthly average

expenditure by twelve:

$$X_j^d = \frac{\sum_i fwt_i^d \sum_t c_{i,j,t}^d}{\sum_i fwt_i^d \; MO\_SCOPE_i^d} \times 12$$

where fwt_i^d is the frequency weight for household i at decile d, $c_{i,j,t}^d$ refers to the annual consumption on category j by household i at decile d and $MO\_SCOPE_i^d$ identify the number of months per year household i reported its expenditures. The decile-level expenditure weight for category d can then be computed as:

$$w_j^d = \frac{X_j^d}{\sum_j X_j^d}$$

D Further robustness checks

As a further robustness check, Figure 15 computes the impulse responses excluding all U.S. recession periods from the analysis (1981M07:1982M11, 1990M07:1991M03, 2001M03:2001M11). The results remain qualitatively unchanged with respect to the baseline specification.

Figure 15: Impulse responses of inflation dispersion (without recession periods)

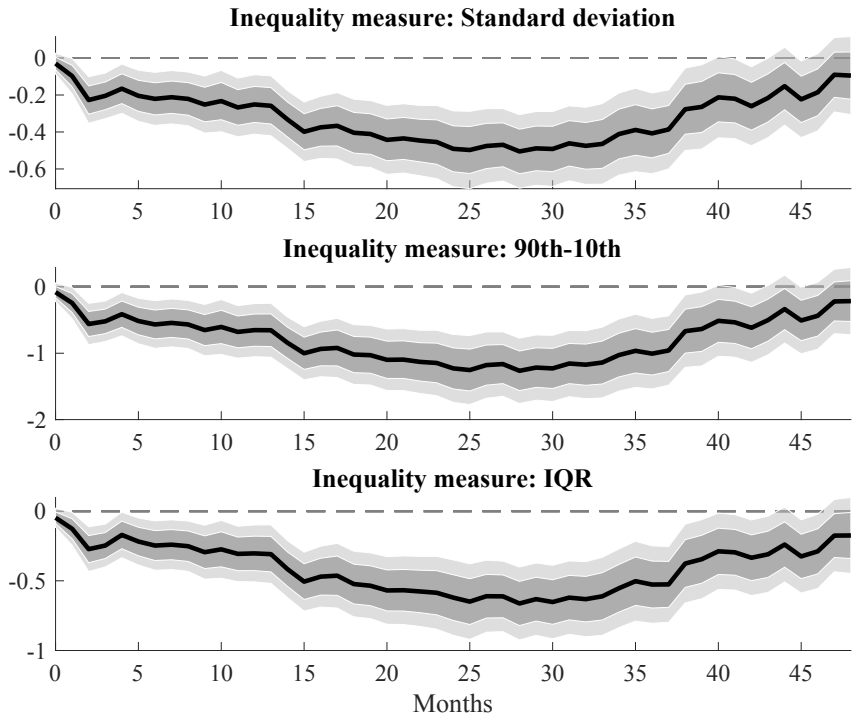

Notes: The figure plots impulse responses to a one percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the respective inflation dispersion measures. The horizontal axis is in months. Dispersion is measured using the cross-sectional standard deviation (top), the difference between the 90th and the 10th percentile of the cross-sectional distribution (middle), and the IQR (bottom). Impulse responses are computed at the monthly frequency using data relative to the period 1980M1:2007M12

As a second set of checks, we assess whether our results are robust to household characteristics. Therefore, following Coibion et al. (2017), the measures of inflation

dispersion are normalized by the number of individuals in the household using the OECD equivalent scale. This variable assigns a value of 1 to the first household member, of 0.7 to each additional adult and of 0.5 to each child. The results are reported in Figure 16.

The magnitude of the impulse responses is comparable to the baseline specification. Moreover, after a monetary policy shocks the decrease in inflation dispersion is more persistent while the confidence intervals are slightly bigger with respect to the baseline case. However, the responses are still statistically significant across the three measures of dispersion.

Moreover, we test whether our results are specific to the shock series we chose (i.e. Romer and Romer, 2004). We confront our estimates with the ones obtained using the shocks proposed by Barakchian and Crowe (2013) and Gertler and Karadi (2015). They identify a monetary policy shock by measuring the change in the expected federal funds target rate in a short time window around the FOMC meetings.

Although the high-frequency approach has been widely used, the two time series are available only starting from the first month of 1990, making them not a viable option for our analysis. In Figure 17 we compare the impulse responses from the local projection using the alternative shocks against our baseline Romer and Romer (2004) shocks. We use the cross-sectional standard deviation as dispersion measure and the data is restricted from 1990M1 to 2007M12 to make the responses comparable.

Removing almost ten years from our sample reduces the magnitude of our baseline response (blue line). However, the coefficients are still largely significant confirming that inflation dispersion decreases after a contractionary monetary policy shock. The coefficients estimated using the Barakchian and Crowe (2013) shocks are close to the zero line for most of the horizon considered but become negative and significant toward the end of the period.

Figure 16: Impulse responses of inflation dispersion (OECD equivalent scale)

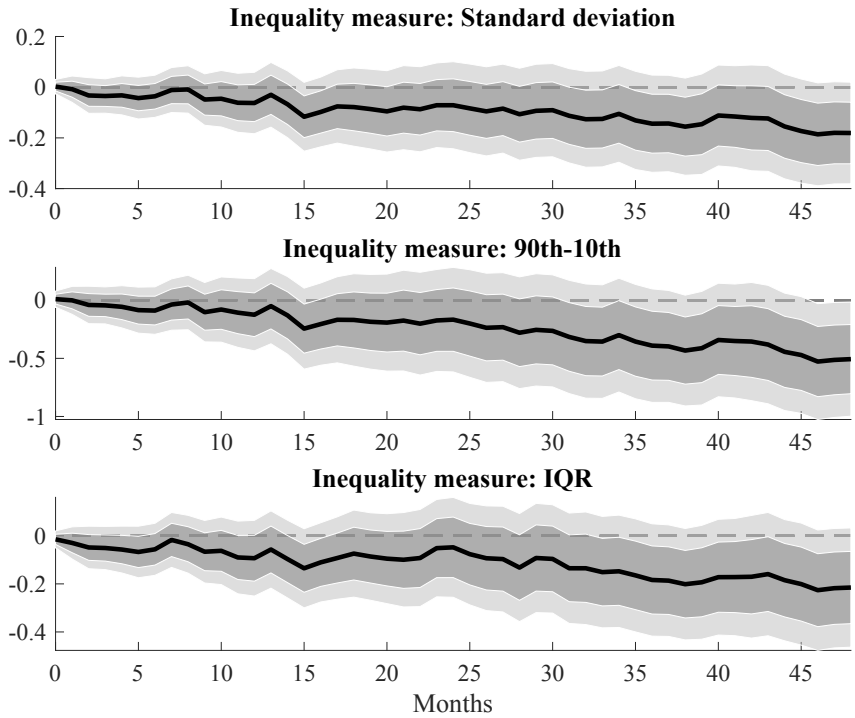

Notes: The figure plots impulse responses to a one percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the respective inflation dispersion measures. Dispersion is measured using the cross-sectional standard deviation (top), the difference between the 90th and the 10th percentiles of the cross-sectional distribution (middle), and the IQR (bottom). The dispersion measures are adjusted using the OECD equivalent scale to control for the household size. Impulse responses are computed at the monthly frequency using data relative to the period 1980M0:2007M12.

The shape and the magnitude of the response computed with Gertler and Karadi (2015) shocks are more in line with our baseline result. Although the plot is noisier than the one obtained with Romer and Romer (2004) shocks, one can notice how it becomes significant only for negative values. Overall, the results from alternative monetary policy shocks do not contradict our main findings and point towards a redistributional role played by monetary policy in terms of inflation dispersion.

Figure 17: Impulse responses of inflation dispersion using alternative shocks

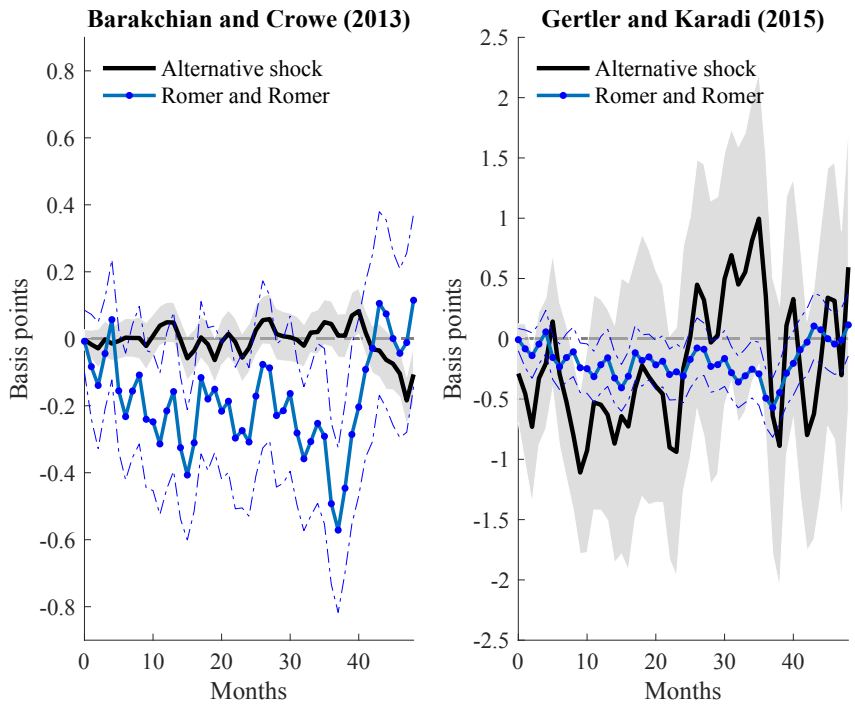

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as the 1.65 standard deviation confidence intervals for the inflation dispersion measured as the cross-sectional standard deviation. The horizontal axis is in months. Impulse responses are computed at the monthly frequency using data for the period 1990M1:2007M12.

In addition, one might be concerned that part of the inflation heterogeneity we measured is driven by differences in consumption patterns across US states rather than along the income distribution. Since the BLS does not provide price indices at the state level but only at the division level (Northeast, Midwest, South and West), we compute the cross-sectional standard deviation of inflation for the four divisions using expenditure weights as well as price indices at division level[23]. As one can

[23] A more limited number of price indices are available at division level therefore we used the following expenditure categories: Food at Home, Food Away from Home, Alcohol, Rented Dwellings, Owned Dwellings, Household Furnishings and Operations, Utility, Apparel, Private Transportation, Public Transportation, Gasoline, Medical, Education and Miscellaneous.

notice in Figure 18, there are no significant differences across divisions in the way our measure of inflation dispersion reacts to a monetary policy shock.

Figure 18: Impulse responses of inflation dispersion across US divisions

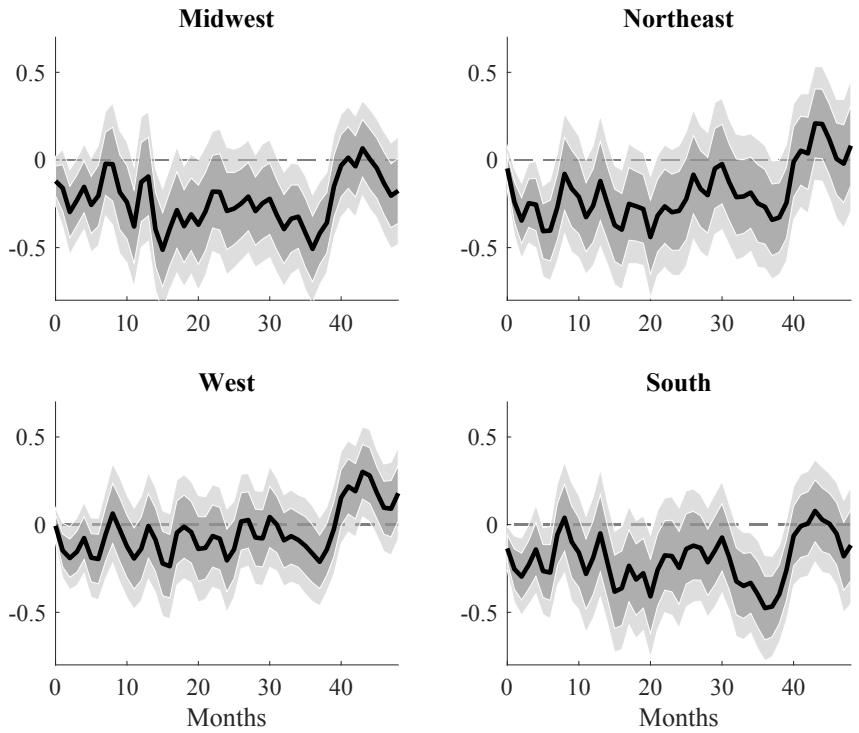

Notes: The figure plots impulse responses to a one percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the inflation dispersion measured as the cross-sectional standard deviation for the four US regions. Impulse responses are computed at the monthly frequency using data relative to the period 1980M1:2007M12

Finally, we follow Ramey (2016) and use as controls the lags of the dependent variable, the shock, industrial production, consumer price index, commodity prices, the policy rate, and unemployment rate in the U.S. Figure 19 reports the result for the cross-sectional standard deviation, which significantly decreases after a monetary policy shock.

Figure 19: Impulse responses of inflation dispersion using Ramey (2016) controls

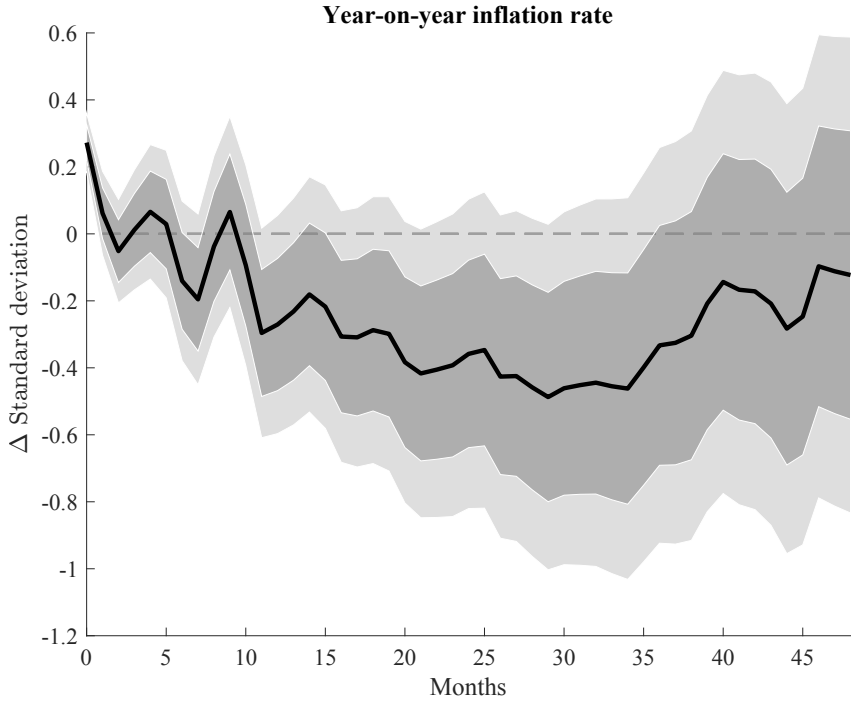

Notes: The figure plots impulse responses to a one percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the inflation dispersion measured as the cross-sectional standard deviation. The controls include the lags of the dependent variable, the shock, industrial production, consumer price index, commodity prices, the policy rate and the unemployment rate. Impulse responses are computed at the monthly frequency using data relative to the period 1980M1:2007M12

E Robustness plots

Figure 20: Impulse responses of inflation dispersion

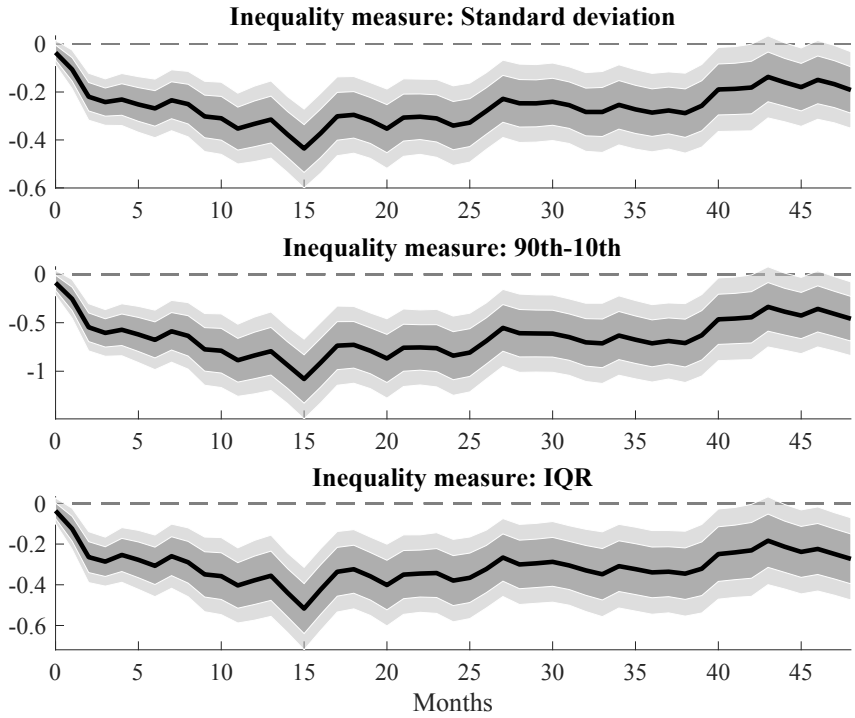

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the respective inflation dispersion measures. The horizontal axis is in months. Dispersion is measured using the cross-sectional standard deviation (top), the difference between the 90th and the 10th percentile of the cross-sectional distribution (middle), and the IQR (bottom). Impulse responses are computed at the monthly frequency using data for the period 1980M1:2007M12.

Figure 21: Impulse responses of inflation dispersion for different lag specifications

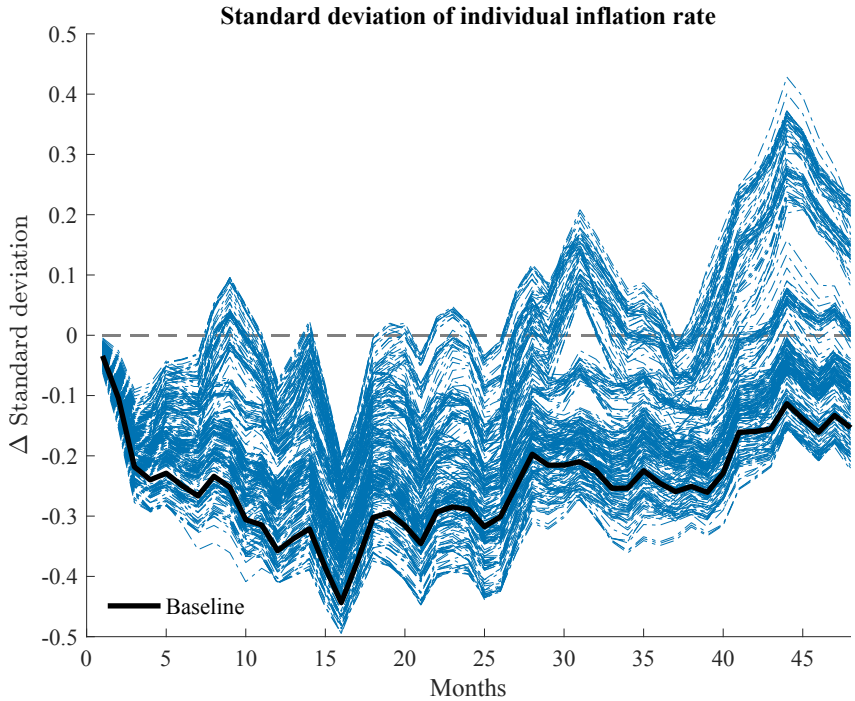

Notes: The figure plots the impulse responses to a one percentage point contractionary monetary policy shock using as the dependent variable the cross-sectional standard deviation and as controls all possible combinations between 5 and 12 lags of the dependent variable and between 24 and 60 lags of the shock variable. The horizontal axis is in months. Impulse responses are computed at the monthly frequency using data relative to the period 1980M1:2007M12.

Figure 22: Impulse responses of the year-on-year inflation rate as well as the median and the standard deviation of the individual inflation rate distribution (LP-IV)

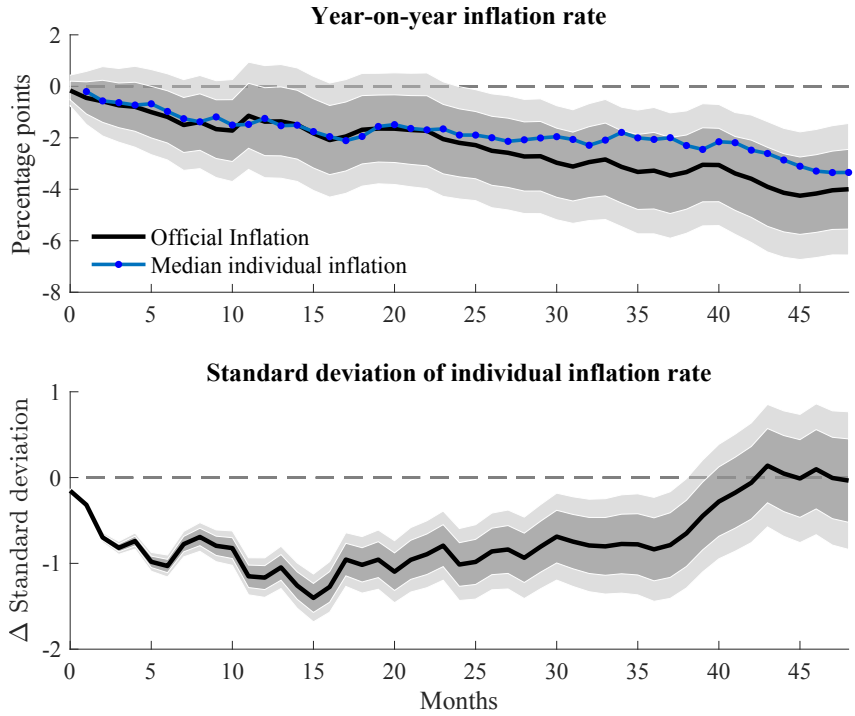

Notes: The figure plots in the top panel the impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the official annual inflation rate (black line) and the median inflation rate (blue line) of the individual inflation rate distribution. The bottom panel reports the impulse response using as the dependent variable the dispersion in inflation measured by the cross-sectional standard deviation. The horizontal axis is in months. Impulse responses are computed at the monthly frequency using data for the period 1980M1:2007M12.

Figure 23: Impulse responses of inflation dispersion (LP-IV)

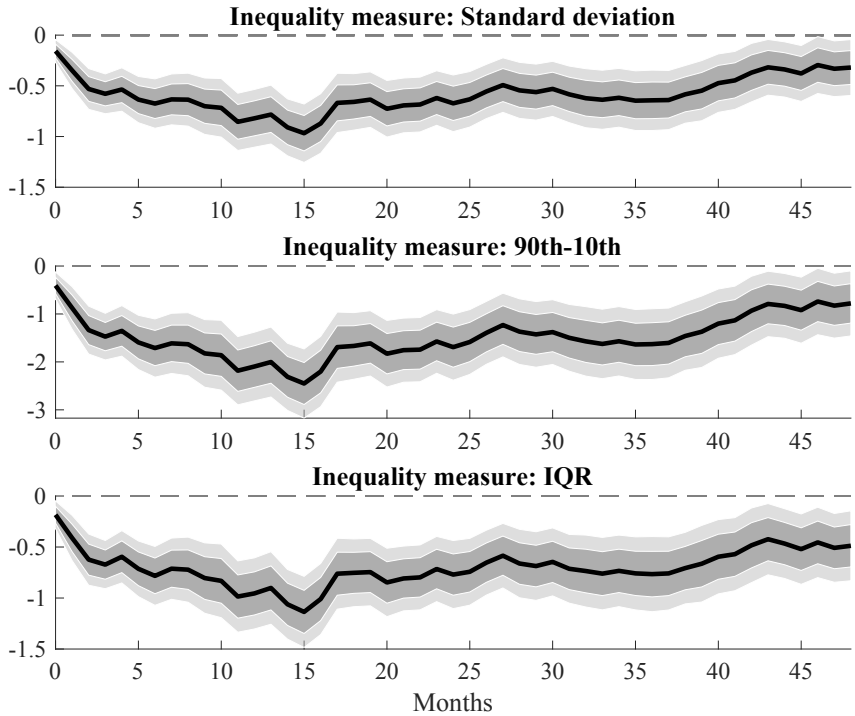

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the respective inflation dispersion measures. The horizontal axis is in months. Dispersion is measured using the cross-sectional standard deviation (top), the difference between the 90th and the 10th percentile of the cross-sectional distribution (middle), and the IQR (bottom). Impulse responses are computed at the monthly frequency using data for the period 1980M1:2007M12.

Figure 24: Forecast error variance decomposition for dispersion measures (LP-IV)

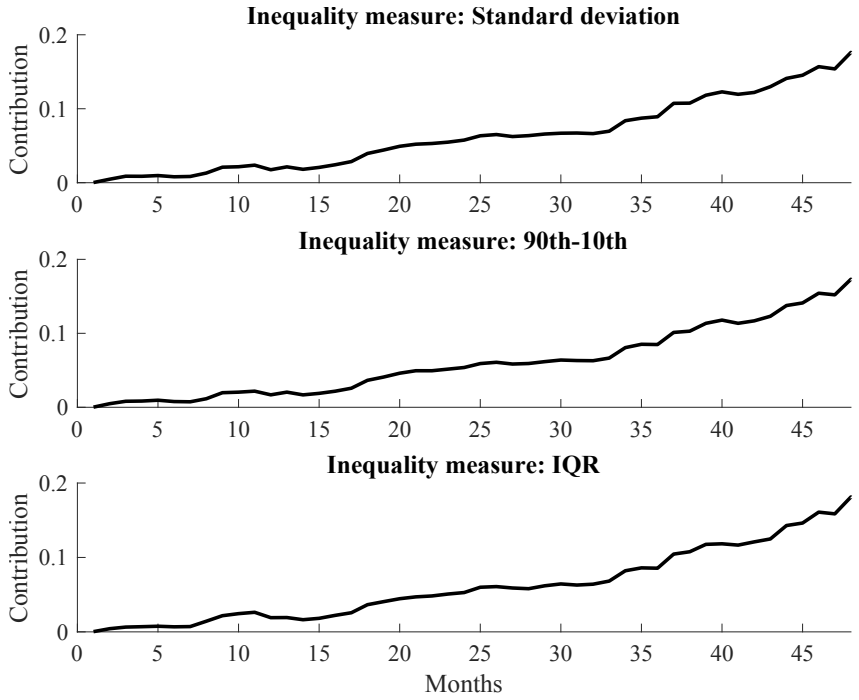

Notes: The figure plots the contribution of monetary policy shocks to the forecast error variance for the respective measure of inflation dispersion at different time horizons (in months).

Figure 25: Impulse responses of inflation dispersion across expenditure, salary, and income deciles (LP-IV)

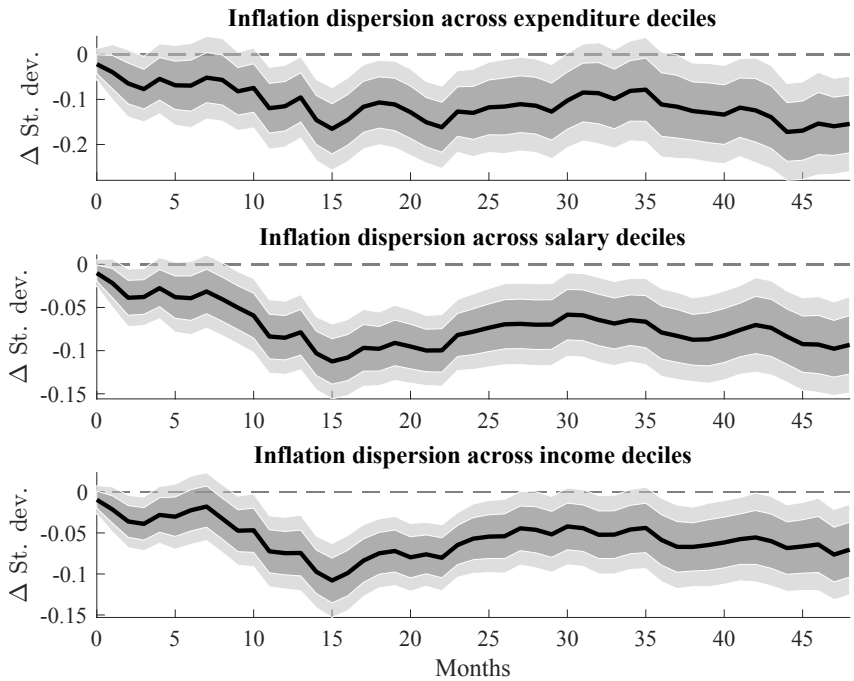

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for inflation dispersion across expenditure (top), salary (middle), and income deciles (bottom). The horizontal axis is in months. Dispersion is measured using the cross-sectional standard deviation. Impulse responses are computed at the monthly frequency using data for the period 1980M1:2007M12.

Figure 26: Impulse responses across expenditure deciles (LP-IV)

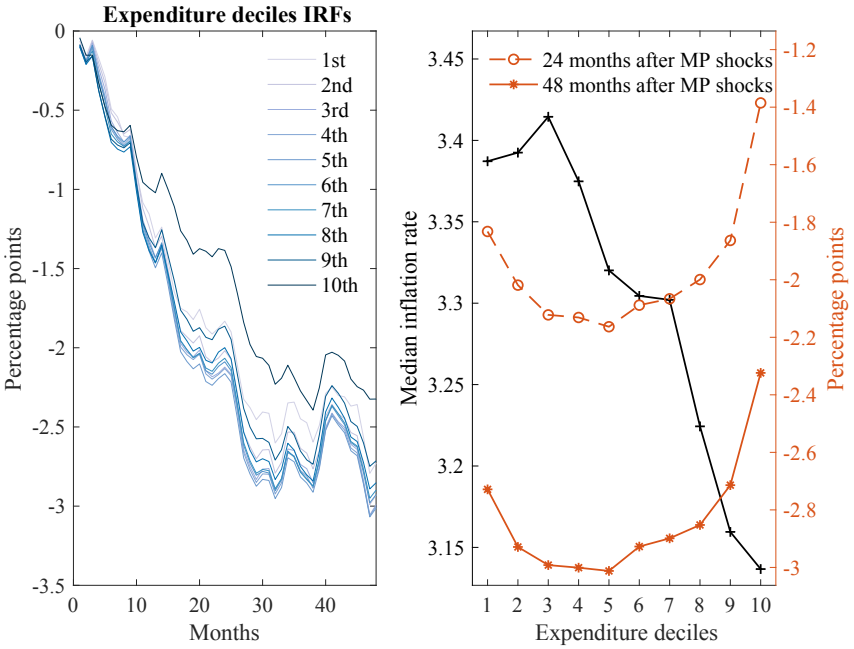

Notes: The left figure plots the impulse responses of the different expenditure deciles following a one percentage point increase in the monetary policy shock, with monthly data for the period 1980M1:2007M12. The right plot reports the median inflation rate across deciles (left axis) as well as the cross-sectional distribution of the responses after 24 and 48 months after the shock (right axis).

Figure 27: Impulse responses across salary deciles (LP-IV)

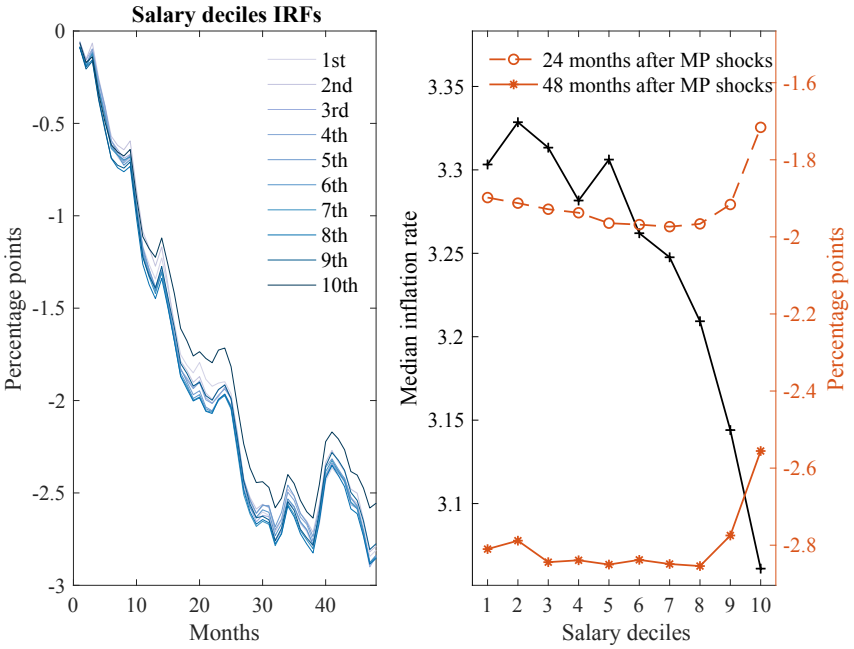

Notes: The left figure plots the impulse responses of the different salary deciles following a one percentage point increase in the monetary policy shock, with monthly data for the period 1980M1:2007M12. The right plot reports the median inflation rate across deciles (left axis) as well as the cross-sectional distribution of the responses after 24 and 48 months after the shock (right axis).

Figure 28: Impulse responses across income deciles (LP-IV)

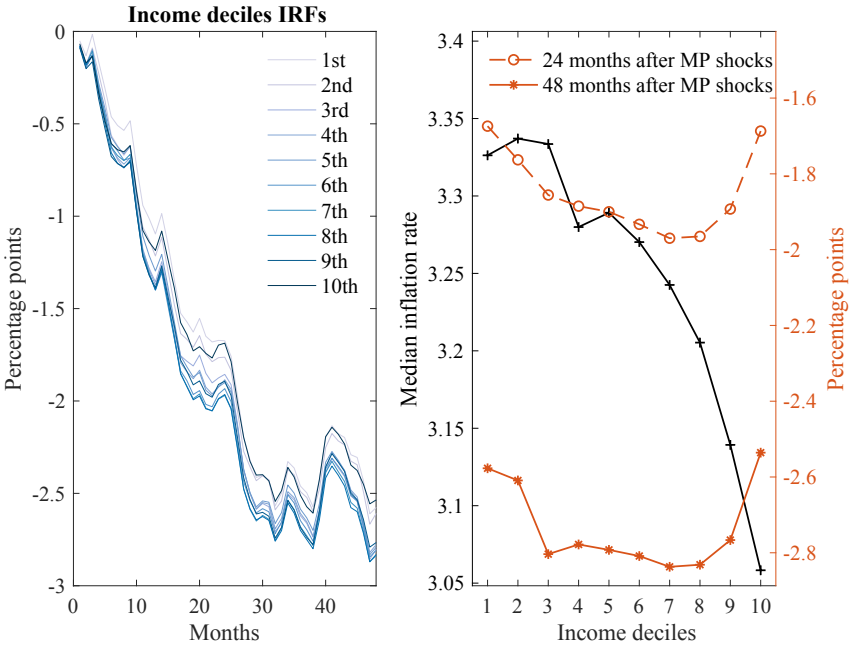

Notes: The left figure plots the impulse responses of the different income deciles following a one percentage point increase in the monetary policy shock, with monthly data for the period 1980M1:2007M12. The right plot reports the median inflation rate across deciles (left axis) as well as the cross-sectional distribution of the responses after 24 and 48 months after the shock (right axis).

Figure 29: Impulse responses of inflation dispersion (without Volcker period)

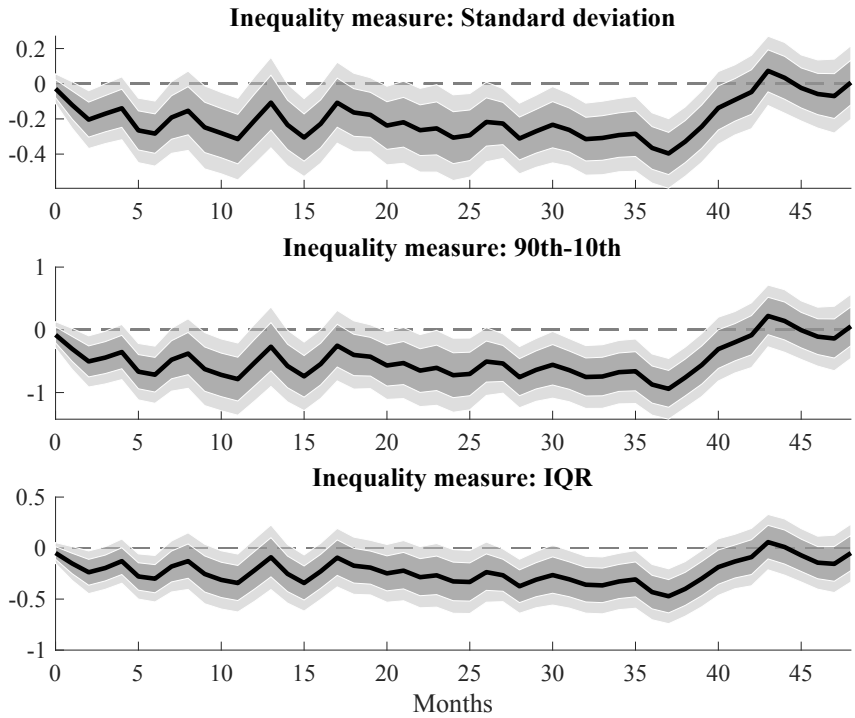

Notes: The figure plots impulse responses to a one percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the respective inflation dispersion measures. The horizontal axis is in months. Dispersion is measured using the cross-sectional standard deviation (top), the difference between the 90th and the 10th percentile of the cross-sectional distribution (middle), and the IQR (bottom). Impulse responses are computed at the monthly frequency using data relative to the period 1985M1:2007M12 in order to exclude the Volcker disinflation period.

Figure 30: Impulse responses of inflation dispersion across expenditure, salary, and income deciles (without Volcker period)

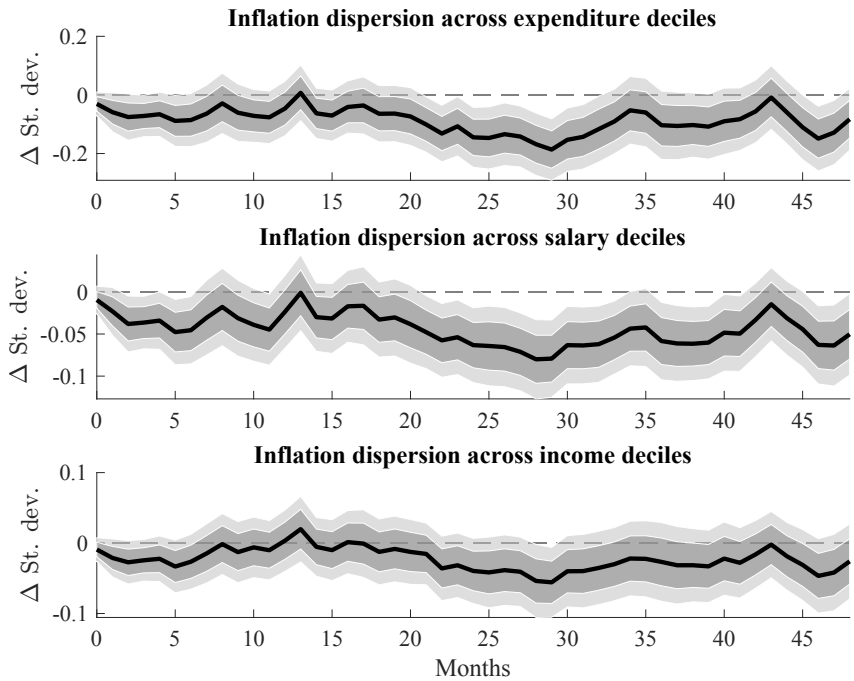

Notes: The figure plots impulse responses to a percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for inflation dispersion across expenditure (top), salary (middle), and income deciles (bottom). The horizontal axis is in months. Dispersion is measured using the cross-sectional standard deviation. Impulse responses are computed at the monthly frequency using data for the period 1985M1:2007M12.

Figure 31: Impulse responses of the cross-sectional standard deviation of inflation (alternative aggregations)

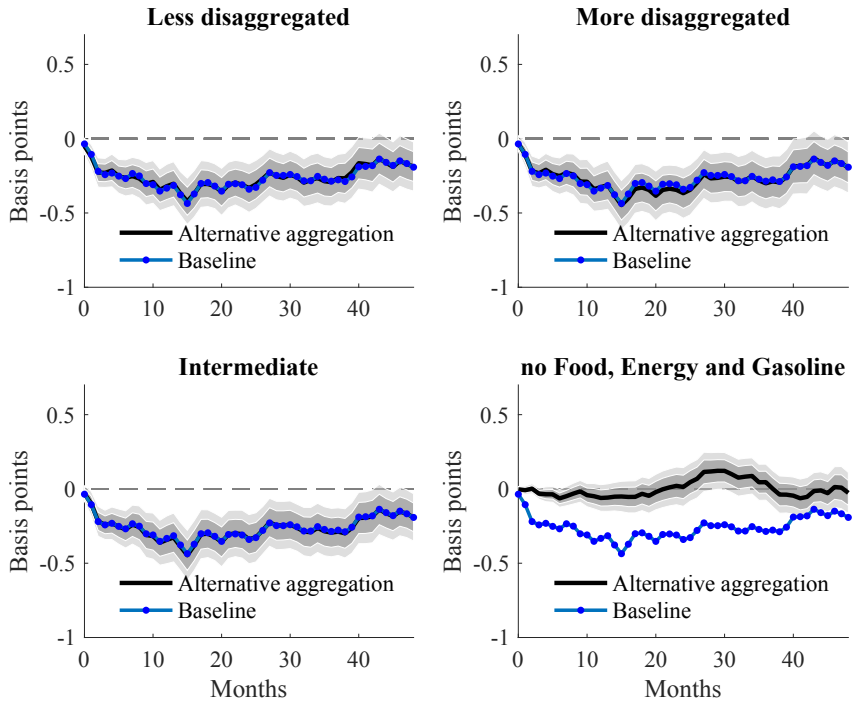

Notes: The figure plots impulse responses of alternatively aggregated inflation rates to a one percentage point contractionary monetary policy shock, as well as one and 1.65 standard deviation confidence intervals for the respective inflation dispersion measures. The solid blue line refers to the impulse response obtained using the baseline categories. Impulse responses are computed at the monthly frequency using data relative to the period 1980M1:2007M12