Price stickiness along the income distribution and the effects of monetary policy*

Javier Cravino
University of Michigan and NBER

Ting Lan
University of Michigan

Andrei A. Levchenko
University of Michigan, NBER, and CEPR

First version: March 2018
This version: October 2018

Abstract

Monetary shocks have distributional consequences if they affect relative prices across goods consumed by different households. We document that the prices of the goods consumed by high-income households are stickier and less volatile than those of the goods consumed by middle-income households. Following a monetary policy shock, the estimated impulse responses of high-income households' consumer price indices are about one-third smaller than those of the middle-income households. We evaluate the implications of these findings in a quantitative multi-sector New-Keynesian model featuring heterogeneous households. The distributional consequences of monetary policy shocks are large and similar to those in the econometric model.

JEL Codes: E31, E52.

Keywords: Inflation, distributional effects, consumption baskets, monetary policy.

*We are grateful to the editor (Yuriy Gorodnichenko), an anonymous referee, Andres Blanco, John Leahy, Indrajit Mitra, and workshop participants at different institutions for helpful comments, to Geoffrey Paulin, David Johnson, Steve Henderson, and John Bieler for help with the CES data, and to Sam Haltenhof for excellent research assistance. Financial support from the National Science Foundation under grant SES-1628879 and from the Washington Center for Equitable Growth is gratefully acknowledged.
Email: jcravino@umich.edu, tinglan@umich.edu, alev@umich.edu.

1 Introduction

There is growing recognition that monetary policy shocks have distributional consequences. An active literature argues that monetary policy can have differential effects across various types of agents: savers vs. borrowers (Doepke and Schneider, 2006), financially constrained vs. unconstrained (Williamson, 2008), or young vs. old (Wong, 2016). In turn, the heterogeneity in the impact of monetary policy across agents can determine its overall effectiveness (Auclert, 2017; Beraja et al., 2017; Kaplan et al., 2018). Coibion et al. (2017) show empirically that monetary contractions increase both income and consumption inequality. In all of these contributions, the distributional consequences of monetary policy arise from its heterogeneous impact on the value of agents' income or wealth.

This paper proposes and quantifies a novel mechanism through which monetary policy shocks have distributional consequences. If the effects of monetary shocks on prices are heterogeneous across types of goods (Boivin et al., 2009), and consumption baskets differ across the income distribution (e.g., Almås, 2012), then shocks will differentially affect the prices faced by households of different incomes. We document that the prices of the goods consumed by high-income households are (i) more sticky and (ii) less volatile than those of the goods consumed by middle-income households. We then use both econometric estimates and a New Keynesian DSGE model to quantify the distributional consequences of monetary policy shocks. Both methodologies indicate that these consequences are large relative to the aggregate impact of monetary policy on prices.

Our analysis uses three main sources of data. The first is the US Consumer Expenditure Survey (CES), from which we obtain expenditure shares across detailed product categories for households at different percentiles of the income distribution. The second is the item-level consumer price data from the BLS, which are the most finely disaggregated consumer prices publicly available for the US. Finally, we employ the measures of price stickiness constructed by Nakamura and Steinsson (2008), who report the frequency of price adjustment (i.e. the probability that a price changes in a particular month) for every detailed product category in the US CPI.

We combine these data to compute the average frequencies of price changes for the baskets of goods purchased by households at each income percentile in the CES. We find systematic differences in the price-stickiness of the consumption baskets of different households. On average, 22% of the goods consumed by households in the middle of the income distribution change prices in a given month. However, the frequency of price changes is 24% lower for the goods consumed by the richest percentile.[1] We also compute income-specific consumer price indices (CPIs), following the procedure that the BLS adopts for computing the aggregate CPI.[2] We show similar differences in the volatility of prices faced by different households: the standard deviation of the CPI of the top percentile is 38% lower than that of the CPI of the middle-income households.

These differences across consumption baskets imply that income-specific CPIs may respond differentially to monetary policy shocks. In particular, the CPIs of high-income households should be less responsive to monetary shocks than the CPIs in the middle of the income distribution. We evaluate this hypothesis both econometrically and quantitatively. We first estimate the impulse responses of income-specific CPIs to monetary policy shocks identified using the narrative approach of Romer and Romer (2004), as extended by Coibion et al. (2017). We compute the impulse responses using the local projections method (Jordà, 2005). Our estimates show that after 36 months, the CPIs of high-income households respond by about one-third less to the same monetary policy shocks than the CPIs of the middle-income households. Thus, the differences in price stickiness and inflation volatility across consumption baskets have the expected impact on the differential responses of households-specific CPIs to monetary policy shocks in the data.

[1] These numbers correspond to frequencies of regular price changes (i.e. excluding sales). The results are similar for the frequency of all price changes (including sales).

[2] When building aggregate consumer price indices, the BLS periodically changes the base year for expenditure weights. In computing income-specific CPIs, we follow the BLS procedure for switching base years after 2004. The information on income is less reliable in the CES prior to 2004, and thus we use 2004 household-specific expenditure weights for CPIs prior to 2004. Using official BLS weights or 2004 aggregate weights produces nearly identical pre-2004 aggregate CPI. See Appendix A.2 for more detail.

We then perform a quantitative assessment using a multi-sector, multi-household model with Calvo-style nominal rigidities. In the model, sectors are heterogeneous with respect to their price stickiness, and households are heterogeneous with respect to their income levels and consumption baskets. We calibrate the model to the observed levels of price stickiness and observed cross-household differences in consumption patterns, and simulate the model's response to a monetary policy shock, paying special attention to how a monetary shock differentially affects households. As expected, high-income households' CPIs respond less to a monetary policy shock than middle-income households' CPIs. The difference is once again quantitatively large: after 12 months, the CPI of the households in the top percentile of the income distribution responds by 13% less than that of the middle-income households. We also show that shifting the distribution of income towards households that consume more sticky goods (i.e. more income inequality) would increase the effectiveness of monetary policy, although this effect is modest for realistic changes in inequality.

Our paper draws on, and contributes to, two literatures. The first is the research agenda on the distributional aspects of monetary policy reviewed above. The second is the literature on the differential responses of prices faced by different consumers following macroeconomic shocks. Cravino and Levchenko (2017) document that after a large devaluation in Mexico, consumption price indices of high-income households increased by far less than consumption price indices of the poor. Argente and Lee (2015) show that in the US Great Recession, prices of groceries and general merchandise items consumed by the poorer households increased by more than those consumed by the richer households, while Jaravel (2017) shows that over the past 15 years, product variety increased the most, and inflation was lowest, for the consumption basket of the high-income households. Kaplan and Schulhofer-Wohl (2017) document substantial cross-sectional dispersion in household inflation rates, while Coibion et al. (2015) study the impact of local economic conditions on the geographical variation in prices paid by consumers. Kim

(2018) shows that low-quality brands change prices more frequently than high-quality brands within narrow product categories, and evaluates the impact of monetary policy across consumers buying goods of different qualities. Ongoing work by Clayton et al. (2018) focuses on differences in price stickiness of goods consumed by, and produced by, college-educated workers. Our paper documents new facts and proposes a novel mechanism that is based on differential price stickiness of consumption items along the income distribution.

The rest of the paper is organized as follows. Section 2 lays out a simple model that illustrates the main mechanism at work, and highlights the key objects of interest that should be the focus of the empirical analysis. Section 3 describes the data and documents consumption basket differences across households. Section 4 presents the econometric evidence, and Section 5 presents the quantitative model and reports the responses of household-specific inflation to an aggregate monetary shock. Section 6 concludes.

2 A simple sticky price model

Before presenting our data, we describe a simplified sticky price model to build intuition on how aggregate shocks can have distributional consequences when nominal rigidities are heterogeneous across goods and households consume different baskets of goods.

Setup: Consider a two-period economy populated by H types of households indexed by h, each consuming a different basket of goods. In the first period, the state of the world is known, and in the second period the economy can experience one of infinitely many shocks or states, s.[3] The (log) price of the consumption basket (i.e. the CPI) consumed by

[3]The set of shocks can include monetary shocks, but at this stage we do not need to specify the exact nature of the shocks.

household h in period t is given by

$$p_t^h(s) \equiv \sum_j \omega_j^h p_{j,t}(s),$$

where ω_j^h is the share of goods from sector j in household h's consumption basket. We define the aggregate price index as $p_t(s) \equiv \sum_h s^h p_t^h(s) = \sum_j \omega_j p_{j,t}(s)$, where s^h denotes household h's share in the aggregate consumption expenditures, and $\omega_j \equiv \sum_h s^h \omega_j^h$ is the economy-wide expenditure share in sector j.

Sectoral goods are aggregates of a continuum of intermediates that are produced by monopolistically competitive firms. We introduce price stickiness by assuming that in the second period, only a fraction θ_j of producers in each sector j can observe the realization of the state before setting their prices. The remaining producers must set prices before observing the realization of the state. To isolate the role of sectoral differences in price rigidities, we assume all producers operate the same CRS technology and set constant markups. In the first period, all the producers know the state and so they set the same price, which we label p_1. In the second period, all producers that observe the state set the same price, which we label $\bar{p}_2(s)$. The producers that don't observe the state set a price that we label p_2^e. Note that p^e is not a function of the state. Without loss of generality we assume that the shocks are mean zero, so that $p_2^e = p_1$.

The average price in sector j in the second period is then given by:

$$p_{j,2}(s) = \theta_j \bar{p}_2(s) + \left[1 - \theta_j\right] p_1. \tag{1}$$

Let $\pi^h \equiv p_2^h(s) - p_1$ define the household-specific inflation rate. The difference in inflation faced by two households, h and h', is:

$$\pi^h(s) - \pi^{h'}(s) = \left[\bar{p}_2(s) - p_1\right] \sum_j \left[\omega_j^h - \omega_j^{h'}\right] \theta_j.$$

This expression highlights that the difference between two households' CPIs is driven by the covariance between the differences in their expenditure shares across sectors, $\omega_j^{h'} - \omega_j^h$, and the price stickiness of those sectors, θ_j. Households that consume less price-sticky goods will experience larger CPI changes following a shock than households consuming relatively more price-sticky goods. Dividing by the aggregate inflation $\pi(s) \equiv p_2(s) - p_1$, yields an expression relating the differences in household-specific inflation to objects that can be measured in the data:

$$\frac{\pi^h(s) - \pi^{h'}(s)}{\pi(s)} = \frac{\bar{\theta}^h - \bar{\theta}^{h'}}{\bar{\theta}}, \tag{2}$$

where $\bar{\theta}^h \equiv \sum_j \omega_j^h \theta_j$ and $\bar{\theta} \equiv \sum_h s^h \bar{\theta}^h$. Note that this expression is independent of the realization of the state.

Discussion: Equation (2) shows how aggregate shocks can have distributional consequences when price rigidities are heterogeneous across goods and households consume different baskets of goods. In this simple model where all firms face the same costs and markups are constant, the weighted average frequencies of price changes, $\bar{\theta}^h$, are sufficient statistics for all the distributional consequences, irrespective of the nature of the aggregate shocks. Equation (2) states that, in response to a shock that generates positive inflation, inflation will be relatively high for households consuming goods with relatively more flexible prices (i.e. high $\bar{\theta}^h$).

To get a sense of the magnitude of these distributional consequences we can do a back of the envelope calculation using US data (described in detail below). In our data, $\bar{\theta}^t \approx 0.17$ for households in the top percentile of the income distribution, $\bar{\theta}^m \approx 0.22$ for households at the middle of the income distribution, and $\bar{\theta} \approx 0.21$. These numbers result in $\frac{\bar{\theta}^t - \bar{\theta}^m}{\bar{\theta}} \approx -0.24$, which indicates that a shock that increases the aggregate CPI relative to its unconditional mean by 1% will also generate a -0.24% gap between the price of the consumption baskets consumed by the top vs. the middle of the income distribution.

The simple model also illustrates the connection between sectoral price stickiness and sectoral price volatility. From (1), we can see that sectoral inflation, $\pi_j(s) \equiv p_{j,2}(s) - p_1$, is less volatile in more sticky-priced sectors:

$$\sigma_{\pi_j} = \theta_j \sigma_{\bar{p}},$$

where σ_{π_j} is the standard deviation of inflation in sector j price, and $\sigma_{\bar{p}}$ is the unconditional standard deviation of $\bar{p}_2(s)$. The ratio of standard deviations of sectoral inflation relative to the standard deviation of aggregate inflation is then given by the ratio of the sectoral to the aggregate frequency of price changes:

$$\frac{\sigma_{\pi_j}}{\sigma_\pi} = \frac{\theta_j}{\bar{\theta}}, \tag{3}$$

Differences in sectoral price volatility translate into differences in household-level CPI volatility. The standard deviation of household-specific inflation, normalized relative to the aggregate is:

$$\frac{\sigma_{\pi^h}}{\sigma_\pi} = \frac{\bar{\theta}^h}{\bar{\theta}}. \tag{4}$$

Households consuming more price-sticky goods experience less volatile price changes. The following section evaluates the relationships (3) and (4) in the data. Of course, these relationships may not hold if the standard deviation of the desired price change $\sigma_{\bar{p}}$ is sector-specific (as would be the case for example if there are sector-specific shocks).

To summarize, our illustrative model establishes that in order to understand how the CPIs of different households react to monetary or other shocks, we must examine the differences in price stickiness of consumption baskets across households. In addition, it suggests a one-to-one relationship between sectoral price stickiness and sectoral price volatility. Thus, a closely related object to be examined in the data is differences in inflation volatility across households.

3 Empirical findings

This section describes our data sources and documents our two empirical findings on how consumption baskets differ across the income distribution. It then evaluates the relationship between frequencies of price changes and inflation volatility suggested by equations (3) and (4). Appendix A describes in detail the construction of expenditure shares from the CES and of the income-specific CPIs.

3.1 Data

We combine data on expenditure shares from the CES with the item-level consumer prices from the BLS and with the frequency of price adjustment data from Nakamura and Steinsson (2008). The CES contains two main modules, the Interview and the Diary. The Interview module collects responses from about 30,000 households annually, and asks households about the purchases they make in all categories, as well as other demographic information. Each household is interviewed for up to 4 consecutive quarters in the Interview module. The Diary module surveys about 10,000 households per year, at weekly frequency. The Diary questionnaire contains detailed questions about daily purchases, such as groceries. All in all, there are questions on 350 distinct expenditure categories in the Interview module, and on 250 distinct grocery and related categories in the Diary module.

The large majority of households do not report expenditures in all possible categories in a given year. In addition, a different set of households is surveyed in the Interview and in the Diary files, so the full consumption profile (both Diary and Interview module expenditures together) of any particular household is never observed. This means that we cannot compute expenditure shares for each household. Rather, we aggregate households into percentiles and work with percentile-level expenditure shares. Each percentile contains about 300 households responding to the Interview questions, and 100 households responding to Diary questions. Appendix Table A1 in Appendix A.1.5 reports the income cutoffs and average incomes in the selected quantiles of the income distribution. It is important to note that income categories in the CES (such as wage income) are subject to top-coding. Nonetheless, there is a great deal of variation in incomes of households present in the CES, with incomes of the top 5% of households an order of magnitude higher than those at the median. Throughout the paper, the percentiles of the income distribution are defined based on income information in the CES rather than any external data source.

We use these data to compute the measures of income-specific frequencies of price changes, price indices, and price volatility defined in Section 2. The average frequencies of price changes, $\bar{\theta}^h = \sum_j \omega_j^h \theta_j$, combine the income-specific expenditure weights ω_j^h from the CES with the product-specific frequencies of price changes θ_j from Nakamura and Steinsson (2008).[4] To compute them, we match CES expenditure categories to the Entry Level Items (ELIs), a basic category in the CPI for which Nakamura and Steinsson (2008) report frequencies. There are a total of 265 ELI categories. In this exercise, we use the expenditure shares from the year 2015, but the results are quite similar for expenditure shares in other years.

We calculate household-specific inflation as $\pi_t^h \equiv p_t^h - p_{t-1}^h$, where the household-specific price indices are given by $p_t^h \equiv \sum_j \omega_{j,\tau}^h p_{j,t}$. The time-varying income-specific expenditure weights $\omega_{j,\tau}^h$ come from the CES, and are updated following the procedure used by the BLS to compute the aggregate CPI. See Appendix A.2 for the complete description of the procedure.[5] The item-level price indices $p_{j,t}$ also come from the BLS. The item level is the finest publicly available level of disaggregation in the US CPI data (the BLS does not report inflation numbers at the ELI level), and is slightly more coarse than ELI, containing 178 distinct expenditure categories starting in 1998. The price data are monthly, for the period 1969-2008, though prior to 1998 the BLS used a different product classifica-

[4] Nakamura and Steinsson (2008) calculate these frequencies as the fraction of prices that change in a given a month, both for all prices, and for regular prices (excluding sales).

[5] Like aggregate measures of inflation, our household-specific price indices are subject to substitution bias. This bias is second-order, and is likely to be negligible for realistic monetary policy shocks. Online Appendix C shows that the differences in inflation rates computed from a Laspeyres vs. a Paasche price index between 1987 and 2016 are an order of magnitude smaller than the inflation rate.

tion. We take 12-month log-differences to obtain annual growth rates. We then compute the standard deviations of those annual growth rates for the price indices at each income level.

3.2 Two facts about consumption basket differences across households

Fact 1: Prices of goods consumed by middle-income households are relatively flexible. Figure 1 presents the scatterplot of the weighted mean frequency of price adjustment, $\bar{\theta}^h$, for households at each of the 20 quantiles of the income distribution in the CES. Thus, each dot corresponds to 5% of households. The solid line through the data is the local polynomial fit, and the shaded area is the 95% confidence interval. The left panel depicts $\bar{\theta}^h$ when θ_j is measured as the frequency of regular (non-sale) price changes, while the right panel measures θ_j as the frequency of all price changes, including sales. Mean frequencies of price changes are hump-shaped along the income distribution: middle-income households consume goods with more frequent price changes, while high- and low-income households consume goods with less frequent price changes.

Figure 1: Weighted mean frequency of price changes

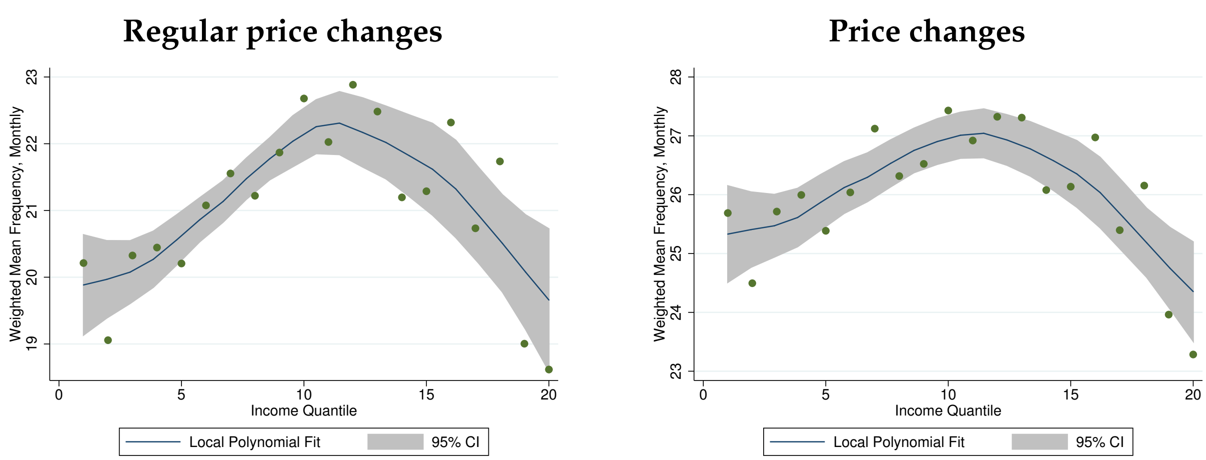

Notes: This figure plots the weighted mean frequency of price changes for households in 20 quantiles of the income distribution. Each dot represents 5% of the income distribution.

Table 1 summarizes the underlying magnitudes. It reports, for different slices of the income distribution, the weighted mean frequency of price adjustment. For the households around the median – the 40-60 income percentiles – the frequency of regular price adjustment is 22.16 percent per month. For all the households from the 1st to the 95th percentile, that frequency is 21.17 percent per month. By contrast, the frequency falls to 19.27 for the households in the 96th to 99th percentile, and further to 16.82 for the top percentile in the distribution. Thus, the weighted mean frequency of price adjustment is some 25% lower for the households in the top 1% of income compared to the households around the median income. Including sales, the results are quite similar. In particular, the top 1% of the income distribution has an 18% lower weighted mean frequency of price adjustment than the middle of the income distribution.

Table 1: Weighted mean frequency of price changes and CPI volatility at different points on the income distribution

	Income percentile			
	40-60	1-95	96-99	100
Frequency of price changes:				
Regular prices	22.16	21.17	19.27	16.82
All prices (incl. sales)	26.90	26.16	23.75	22.17
Standard deviation of CPI:	0.021	0.020	0.015	0.013

Note: This table reports the weighted mean frequency of price changes, and the standard deviation of the 12-month log change in CPI for consumers of different incomes.

Fact 2: Prices of goods consumed by middle-income households are relatively volatile. Figure 2 reports the standard deviation of π_t^h, the income-specific inflation. Inflation volatility is also hump-shaped along the income distribution. The households with middle incomes experience the highest inflation volatility, whereas the lowest volatility is found at the top of the income distribution. The bottom of Table 1 reports the values of

Figure 2: Standard deviation of the changes in consumption price indices

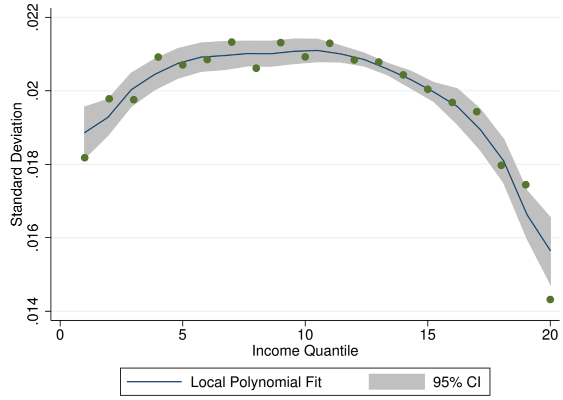

Notes: This figure plots the standard deviation of the 12-month log difference in the consumption price indices for households in 20 quantiles of the income distribution. Each dot represents 5% of the income distribution.

the standard deviation of inflation faced by consumers of different incomes. The annual inflation rate has a standard deviation of 0.020 for consumers in the bottom 95% of the income distribution, and 0.021 for consumers in the middle (40-60th percentiles). By contrast, the standard deviation of annual inflation is 0.015 for households in the 96th to 99th percentile of the income distribution, and 0.013 for those in the top 1%.

Discussion: What consumption patterns are responsible for these differences in price stickiness and volatility across baskets? Online Appendix Table A3 reports the 10 consumption items with the largest differences in the expenditure shares between the middle 20% of the income distribution and the top 1% of the income distribution.

The top categories in which the middle-income consumers exhibit highest expenditure shares relative to the top 1% are mainly goods such as Gasoline, Electricity, Motor Vehicle Insurance, and Used Cars. The items with the largest expenditure shares of the top 1% relative to the middle income consumers are mostly services, such as Elementary School and College Tuition, Child Care, Airfare, Domestic Services, and Club Membership Fees. Among the 10 categories consumed more intensively at the middle of the income distribution, the frequency of monthly price adjustment is in excess of 30%. Among the 10 items most disproportionately consumed by the top 1%, the frequency of regular price adjustment is 16%, and total price adjustment 18%. In either case, the difference in average price adjustment frequency between these two sets of items is pronounced. The pattern of price stickiness is not universal. Among the top 1%'s (relative) top 10 items is Airfare, with price adjustment frequency of almost 60% per month. On the flip side, General Medical Practice and Limited Service Meals are in the middle 20%'s top 10, and among the price-stickiest categories.

The left panel of Figure 3 plots the frequency of the regular price adjustment on the y-axis against the difference in the expenditure shares between the top 1% and the middle 20%, with positive values meaning that the top 1% has higher expenditure shares in that category. The majority of categories are concentrated on 0, implying that the high- and the middle-income categories have similar expenditure shares. There is a large range, however, and all in all the relationship between these relative shares and the frequency of price adjustment is negative. The correlation between the x-axis and y-axis variables is -0.251.

Figure 3: Expenditure differences, frequency of price changes, and standard deviation of price changes

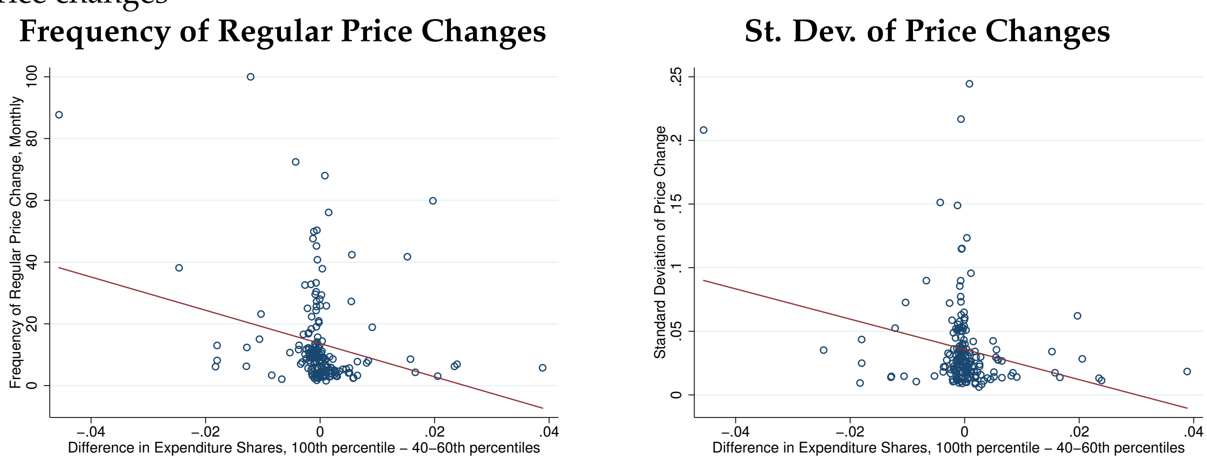

Notes: The left panel plots the frequency of price changes against the difference in sectoral expenditure shares between households in the top 1% and the middle 20% of the income distribution. The right panel plots the standard deviation of 12-month log price change against the difference in sectoral expenditure shares between households in the top 1% and the middle 20% of the income distribution. Both panels include the OLS fit through the data.

The categories with the largest expenditure share differences also differ substantially in the standard deviation of item-level price changes. The mean standard deviation of 12-month log price changes in the set of goods consumed most disproportionately by the middle-income households is 0.049, more than double the 0.023 mean in the set of goods consumed by high-income households.

The outlier sector here is Gasoline, whose standard deviation is 0.208, and which is also the sector with the single largest expenditure share discrepancy – in either direction – between the middle- and high-income households. But the differences persist even if we focus on the median standard deviation, or drop Gasoline when computing the mean.[6] The right panel of Figure 3 displays the scatterplot of the standard deviation of log price change at the item level against the expenditure share difference between the high- and middle-income consumers. Once again, most expenditure share differences are close to zero. Nonetheless, the correlation between the expenditure share differences and standard deviation of price changes is negative at -0.255.

3.3 Frequency of price changes and inflation volatility

This section evaluates the relationship between frequency of price changes and inflation volatility suggested by equations (3) and (4) of Section 2, by providing the data counterparts of those postulated relationships. The left panel of Figure 4 plots the empirical counterpart of (3), along with a 45-degree line. As (3) expresses both the right- and left-hand side variables relative to the average, we rescale both the product-level standard deviation and the frequency of price adjustment by their means across items. Each dot represents one of the 178 disaggregated CPI items. A positive relationship with a slope close to unity is evident in this plot; the correlation coefficient between these two variables

[6] If we exclude Gasoline, the differences across households reported in Figure 1 and Table 1 are attenuated, but the basic patterns hold. Gasoline appears to be responsible for about half of the difference in the weighted average frequency of price adjustment between the top-income and the middle-income households. Dropping Gasoline, Figure 2 is somewhat modified. It is still true that high-income households have lower inflation volatility than middle-income ones, but now the highest inflation volatility occurs in the bottom income tercile.

is 0.715.

The right panel plots the empirical counterpart of (4), once again with both y- and x-axis variables rescaled by their respective means and adding a 45-degree line. Each dot represents 5% of the income distribution, as in Figures 1-2. There is an evident positive relationship between these two variables, with the correlation coefficient of 0.643. Households consuming more flexible-priced goods tend to experience higher CPI volatility. This is not surprising, as we are in effect plotting the y-axes of Figures 1 and 2 against each other, and both follow a similar inverse U-shape with respect to income quantile.

Figure 4 is consistent with the prediction of our time-dependent pricing model from Section 2 that more sticky sectors should have less volatile prices. If the frequencies of price adjustment are endogenous, it could be that the differences in the volatility of sectoral shocks are what drives the frequency of price adjustment. Note that irrespective of the direction of causality, the correlation between inflation volatility and frequency of price adjustment across households of different incomes is that depicted in Figure 4. The following section shows that some households are more sensitive to monetary policy shocks than others. For those results, we do not need to specify whether the difference in the frequency of price adjustment across sectors is exogenous or driven by the volatility of sectoral shocks.[7]

4 Impulse responses of income-specific CPI to monetary policy shocks

The previous section shows that prices of goods consumed by high-income households are more sticky and less volatile than those of the goods consumed by middle-income households. This suggests that monetary shocks can have distributional consequences by affecting the relative prices of consumption baskets of households at different points

[7]Indeed, Boivin et al. (2009) provide evidence that prices of more volatile goods react systematically more strongly to monetary policy shocks.

Figure 4: Stickiness and volatility

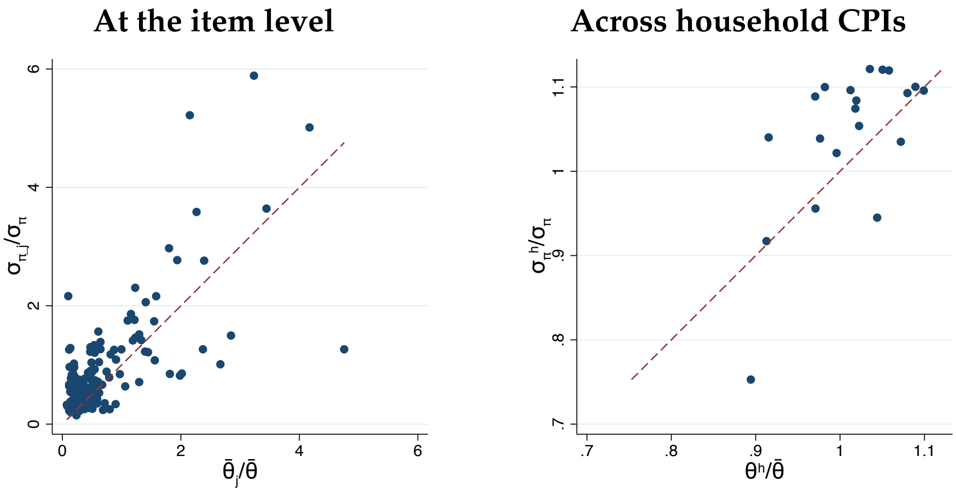

Notes: The left panel plots the standard deviation of 12-month log price change at the item level vs. the frequency of price adjustment for that item. The right panel plots the standard deviation in the 12-month log change in overall household CPI against the weighed mean frequency of price adjustment for that household type; each dot represents 5% of the income distribution. Both plots include the 45-degree line.

on the income distribution. We now present evidence that monetary policy shocks indeed lead to smaller CPI changes for households at the top of the income distribution relative to the middle. Our baseline specification adopts the local projection method of Jordà (2005) to estimate the responses of income-specific CPIs to monetary policy shocks. Online Appendix B presents impulse responses of income-specific price indices using the FAVAR methodology following Bernanke et al. (2005) and Boivin et al. (2009).

The local projection method estimates regressions of the dependent variable at horizon $t+s$ on the shock in period t and uses the coefficient on the shock as the impulse response estimate. We estimate the following series of regressions:

$$p_{t+s}^{h} - p_{t}^{h} = \alpha_s + \theta_s shock_t^{RR} + \sum_{j=1}^{J} \beta_{s,j}(p_{t+1-j}^{h} - p_{t-j}^{h}) + \sum_{i=1}^{I} \gamma_{s,i} shock_{t-i}^{RR} + \epsilon_{t+s}. \tag{5}$$

Here, p_t^h is the log of income-specific CPIs, and $shock^{RR}$ is the Romer and Romer (2004) narrative-based measure of monetary policy shocks from Coibion et al. (2017). The control variables include 48 lags of the shocks ($I = 48$) and 6 lags of monthly income-specific inflation ($J = 6$). The coefficient θ_s gives the response of income-specific prices at $t+s$ to a monetary policy shock at t. We estimate impulse responses over a horizon of 48 month with $s = 0, 1, \ldots, 48$.

In our application, we estimate the impulse response of income-specific prices for each income percentile. We use monthly data for the sample period 1969m1 to 2008m12. Figure 5 plots the estimated impulse responses of income-specific prices for selected percentiles to a 100-basis-point of contractionary monetary policy shock. The consumer price indices of the high-income households react substantially less to monetary policy shocks than those for the middle of the income distribution. The difference is economically meaningful. After 36 months, the top-1% households' CPI responds by 38% less, and the 96-99th percentile households by 26% less, than the CPI of the households in the middle of the income distribution (40-60th percentiles). After 48 months, the differences are still 33% and 22%, respectively.

Figure 5: Income-specific CPI impulse responses to a monetary policy shock

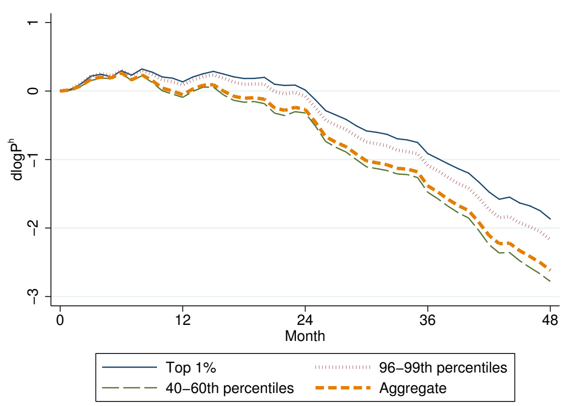

Notes: This figure plots the impulse responses of income-specific price indices to a monetary policy shock identified using the narrative approach of Romer and Romer (2004), as extended by Coibion et al. (2017). The impulse responses are computed using the local projections method (Jordà, 2005)

Our main object of interest is not the overall response of prices to a monetary shock, but rather the differential response of the CPIs of different households. We estimate a version of equation (5) using the difference between the (log) CPI of the top 1% and the (log) CPI of the middle 20% of the income distribution, and the difference between the top 1% and the aggregate CPI to quantify the effect of monetary policy on inflation faced by different households. Figure 6 plots the difference in the response of the CPIs of different households. The dark and light grey areas indicate 1 and 1.65 standard deviation confidence intervals, respectively. The figures show that following a contractionary monetary shock, the price level for the the top income households falls by less than the price level for middle income households (so that the difference between the two is positive). The difference is statistically significant, and is about a third of the size of the response of aggregate inflation to the same monetary shock reported in Figure 5.

The econometric evidence thus suggests that monetary shocks can have large distributional effects across households of different incomes. The following section complements this evidence using a New Keynesian model that quantifies the mechanisms described in Section 2.

Figure 6: Differences in inflation changes between income groups

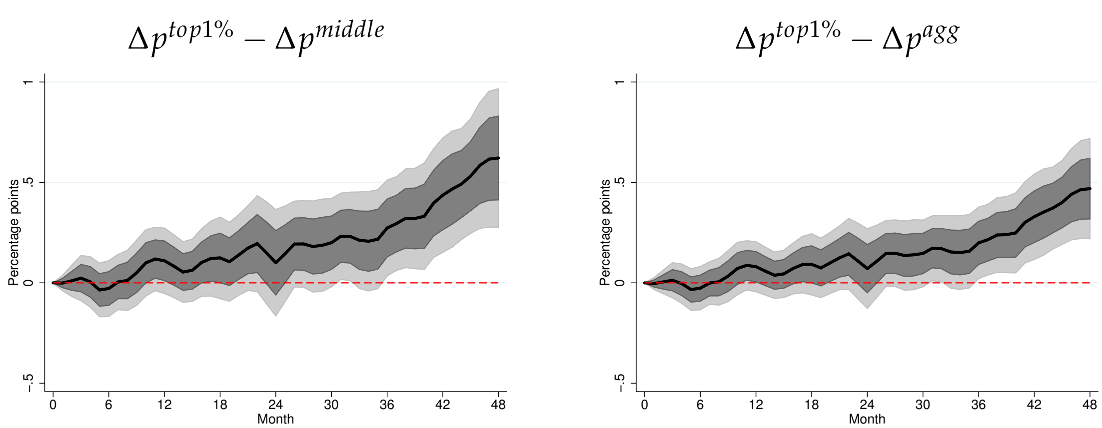

5 Quantitative framework

This section sets up a sticky price model with multiple households and sectors to evaluate how monetary shocks affect consumption price indices for households at different points of the income distribution.

5.1 Setup

Preliminaries: We consider an economy populated by H types of households indexed by h. Households get utility from consuming a bundle of goods produced by J different sectors of the economy indexed by j. Sectoral goods are produced by aggregating the output of a continuum of monopolistic intermediate producers indexed by i. The monetary authority sets the nominal interest rate following a Taylor rule.

Households: Each type of household h has preferences given by:

$$U^h = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[lnC_t^h - N_t^h \right], \tag{6}$$

and faces the budget constraint:

$$P_t^h C_t^h + \Theta_{t,t+1} B_{t+1}^h = W_t A^h N_t^h + T_t^h + B_t^h. \tag{7}$$

Here, C_t^h is the bundle of goods consumed by households of type h, and P_t^h is the price of this bundle. N_t^h and A^h respectively denote labor supply and the efficiency of household h, and W_t is the nominal wage per efficiency unit. B_{t+1}^h is a bond that pays one dollar in $t+1$, and $\Theta_{t,t+1}$ is the date t price of that bond. Finally, T_t^h are transfers to the households from the government and from firms' profits.

The bundle of goods consumed by each type of household is:

$$C_t^h = \left[\sum_j^J \left[\overline{\omega}_j^h \right]^{\frac{1}{\eta}} \left[C_{j,t}^h \right]^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \tag{8}$$

where $C_{j,t}^h$ denotes household h's consumption of final goods from sector j, and $\overline{\omega}_j^h$ is a household-specific taste shifter for sector j. Note that the parameters $\overline{\omega}_j^h$ are associated with a particular household h that has efficiency A^h. These parameters allow us to capture in a reduced form the non-homotheticities that may lead to the cross-household differences in expenditure shares observed in the data. The model will be calibrated directly to household-specific expenditure shares. The price index associated with this bundle is:

$$P_t^h = \left[\sum_j^J \overline{\omega}_j^h P_{j,t}^{1-\eta} \right]^{\frac{1}{1-\eta}},$$

where $P_{j,t}$ is the price of the sector j aggregate. Note that both C_t^h and P_t^h are indexed by h, as the bundle (8) differs across households. Monetary shocks can differentially affect households if households put different weights across sectors and shocks have heterogeneous effects across sectoral prices $P_{j,t}$.

Sectoral demands and technologies: The demand function associated with the bundle (8) is given by:

$$C_{j,t}^h = \overline{\omega}_j^h \left[\frac{P_{j,t}}{P_t^h} \right]^{-\eta} C_t^h.$$

Adding across households, aggregate demand for the final good produced in sector j is

$$P_{j,t} C_{j,t} = \omega_{j,t} \left[\frac{P_{j,t}}{P_t} \right]^{1-\eta} P_t C_t, \tag{9}$$

where $P_t C_t$ are aggregate nominal expenditures, $\omega_{j,t} \equiv \sum_h \overline{\omega}_j^h s^h \frac{\left[P_t^h\right]^{\eta-1}}{\sum_h s_h \left[P_t^h\right]^{\eta-1}}$, and $P_t \equiv \left[\sum_j \omega_{j,t} P_{j,t}^{1-\eta} \right]^{\frac{1}{1-\eta}}$. In these expressions, s^h is the share of household h in aggregate expenditures.

Sectoral goods are produced by aggregating the output of a continuum of intermediate producers according to

$$Y_{j,t} = \left[\int Y_{j,t}(i)^{\frac{\gamma-1}{\gamma}} di \right]^{\frac{\gamma}{\gamma-1}}.$$

Total demand faced by intermediate producer i is then:

$$Y_{j,t}(i) = \left[\frac{P_{j,t}(i)}{P_{j,t}}\right]^{-\gamma} Y_{j,t}. \tag{10}$$

Intermediate good producers: Intermediate producers behave as monopolistic competitors and set prices as in Calvo (1983). The probability that a producer can change its price in any period depends on the sector in which it operates, and is given by θ_j. The producers operate a linear technology

$$Y_{j,t}(i) = \bar{N}_{j,t}(i), \tag{11}$$

where $\bar{N}_{j,t}(i)$ denotes the efficiency units of labor used by producer i. The profit-maximizing price for an intermediate producer that gets to adjust prices satisfies:

$$\bar{P}_{j,t} = \arg\max \left\{ \sum_{k=0}^{\infty} \left[1-\theta_j\right]^k \mathbb{E}_t \left\{ \Theta_{t,t+k} \left[\bar{P}_{j,t} - W_{t+k}\right] Y_{j,t+k}(i) \right\} \right\} \tag{12}$$
$$\text{subject to} \quad (10).$$

Monetary policy: The monetary authority sets nominal interest rates according to a Taylor rule:

$$exp(i_t) = exp(\rho_i i_{t-1}) \left[\Pi_t^{\phi_\pi} \left[\frac{Y_t}{\bar{Y}}\right]^{\phi_y}\right]^{1-\rho_i} exp(\nu_t),$$

where $i_t \equiv -logQ_{t,t+1}$ is the nominal interest rate, $\Pi_t \equiv P_t/P_{t-1}$ is aggregate inflation, and $\bar{Y}$ is the efficient level of output. Finally, ν_t is a monetary shock that satisfies

$$\nu_t = \rho_\nu \nu_t + \varepsilon_{\nu,t}, \tag{13}$$

with $\varepsilon_{\nu,t} \sim N(0, \sigma_{\varepsilon_\nu})$.

Equilibrium: An equilibrium for this economy is a set of allocations for the households $\{C_t^h, C_{j,t}^h, N_t^h\}_{\forall j,h,t}$, sectoral good producers $\{Y_t^j, \left\{Y_t^j(i)\right\}_i, \left\{\bar{N}_t^j(i)\right\}_i\}_{\forall j,t}$, and price policy functions for intermediate producers $\{\bar{P}_{j,t}\}_{\forall j,t}$, such that given prices: (i) households maximize (6) subject to (7); (ii) sector j final producers minimize costs according to equations (9) and (10); (iii) intermediate producers maximize profits by solving (12); and (iv) goods and labor markets clear, $\sum_h C_{j,t}^h = Y_t^j$ and $\sum_h A^h N_t^h = \sum_j A^h \int \bar{N}_{j,t}(i)\, di$.

We now characterize the equilibrium of a log-linearized version of this economy around a non-stochastic steady state, following the tradition in the New Keynesian literature. In what follows, we use lower-case letters to denote the log-deviations of a variable from its non-stochastic steady state. The optimality conditions associated with the household problem are the labor-leisure condition:

$$P_t^h C_t^h = A^h W_t,$$

and the Euler equation:

$$\Theta_{t,t+1} = \beta \mathbb{E}_t \left\{ \frac{P_t^h C_t^h}{P_{t+1}^h C_{t+1}^h} \right\}.$$

Adding the labor-leisure condition across households we obtain that each type of household gets a constant share of nominal consumption expenditures, $s^h \equiv \frac{P_t^h C_t^h}{P_t C_t} = \frac{A^h}{A}$, where $A \equiv \sum_h A^h$. Substituting into the optimality conditions and log-linearizing we obtain:

$$w_t - p_t = c_t, \tag{14}$$

and

$$c_t = \mathbb{E}_t\{c_{t+1}\} - [i_t - \mathbb{E}_t\{\pi_{t+1}\} - \rho], \tag{15}$$

with $\rho \equiv -log\beta$. Goods market clearing implies $y_t = c_t$. Substituting into equation (15) we obtain:

$$y_t = \mathbb{E}_t \{y_{t+1}\} - [i_t - \mathbb{E}_t \{\pi_{t+1}\} - \rho] . \tag{16}$$

The optimal log-price that solves (12) can be written recursively as:

$$\bar{p}_{j,t} = \left[1 - \beta \left[1 - \theta_j\right]\right] w_t + \beta \left[1 - \theta_j\right] \mathbb{E}_t \left[\bar{p}_{j,t+1}\right],$$

and the law of motion for the sectoral price indices is

$$p_{j,t} = \theta_j \bar{p}_{j,t} + \left[1 - \theta_j\right] p_{j,t-1}.$$

Combining we these two equations we obtain a sectoral Phillips curve,

$$\pi_{j,t} = \lambda_j \left[w_t - p_{j,t}\right] + \beta \mathbb{E}_t \left\{\pi_{j,t+1}\right\}, \tag{17}$$

with $\lambda_j \equiv \frac{\theta_j\left[1-\beta\left(1-\theta_j\right)\right]}{\left[1-\theta_j\right]}$. Finally, the Taylor rule is:

$$i_t = \rho_i i_{t-1} + [1 - \rho_i] \left[\rho + \phi_\pi \pi_t + \phi_y \tilde{y}_t\right] + \nu_t. \tag{18}$$

Equations (14)-(18) can be used to solve for all sectoral inflation rates, along with the output gap, real marginal costs, real wages, the nominal interest rate, and the aggregate inflation rate. Sectoral inflation rates can then be used to compute household-specific inflation according to:

$$\pi_t^h = \sum_j \omega_j^h \pi_{j,t}.$$

Note that as a result of taking the first-order approximation, we dropped the time subscripts on the expenditure weights ω_j^h, since changes in prices only affect expenditure shares to a second order. This second-order substitution bias is likely to be negligible for realistic monetary policy shocks, as discussed in Appendix C.

In what follows, we will use the model to ask two questions: (i) what is the effect of a monetary policy shock $\varepsilon_{\nu,t}$ on household-specific inflation?, and (ii) how do changes in the distribution of income s^h affect the response of inflation π_t and the output y_t to a monetary shock?

5.2 Results

5.2.1 Calibration

To evaluate the impact of monetary shocks, we need to assign values for the discount factor β, the coefficients in the Taylor rule, ρ_i, ϕ_π and ϕ_y, the process for the shocks, ρ_ν and σ_{ε_ν}, the sectoral frequencies of price changes, θ_j, $\forall_j$, the sectoral household-specific expenditure shares, ω_j^h, and the household shares in aggregate consumption spending, s^h. We calibrate the model to monthly data and use values for most of these parameters that are standard in the literature. In particular, we set $\beta = 0.96^{1/12}$, which corresponds to an annualized real interest rate of 4 percent, and take the Taylor rule parameters $\rho_i = 0.95$, $\phi_\pi = 1.5$ and $\phi_y = 0.5/12$ and set the persistence of the shocks to $\rho_\nu = 0$, as in Christiano et al. (2010). Finally, we calibrate the model to 265 sectors and 20 household types, and calibrate the frequencies of price changes θ_j and the expenditure shares ω_j^h and s^h using the data from Nakamura and Steinsson (2008) and the CES data presented in Section 3.

5.2.2 Distributional consequences of monetary shocks

We now evaluate the distributional consequences of a monetary shock in this model. Figure 7 plots the impulse response of the household-specific price indices to a one standard deviation shock to $\varepsilon_{\nu.t}$. The figure shows that the shock has distributional effects: prices of the middle-income households are the most sensitive to the shock, and prices are the least sensitive for the top-income households. This is not surprising, since in our model, as in the data, households at the top of the income distribution consume the goods that are the most sticky and thus respond more sluggishly to shocks.

Table 2 reports the price indices faced by households at different points of the income distribution following the monetary shock, expressed relative to the aggregate price index. The table shows that the cumulative response after 6 months of the prices faced by the top 1 percent is about 13% smaller than that of the aggregate price index, and almost 20% smaller than the response of the prices faced by the households at the middle 5 percent of the income distribution. These differences are quite persistent, the cumulative change in prices faced by the richest 1% is still 10% smaller than that faced by the middle income households 18 months after the shock.

Figure 7: Impulse responses of household-specific CPIs to a monetary shock

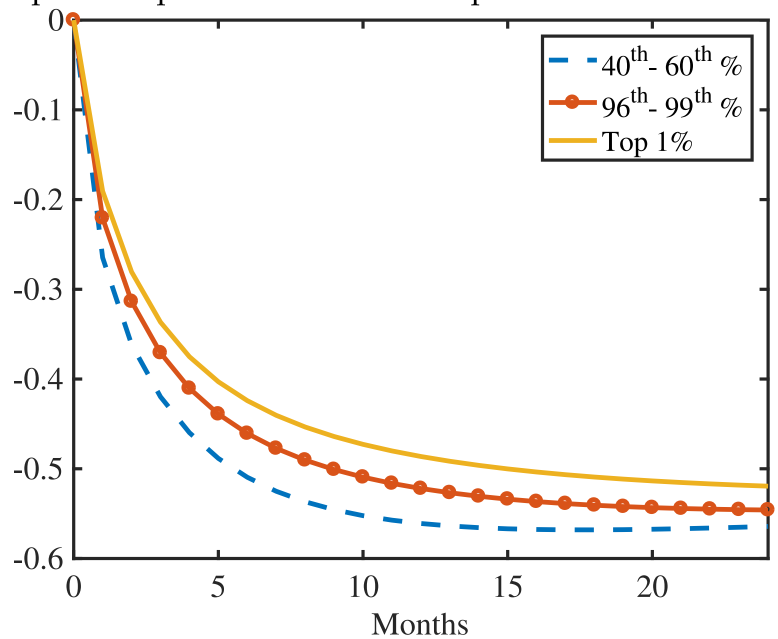

Notes: This figure plots the impulse responses of income-specific CPIs to a monetary policy shock, simulated using the model in this section.

5.2.3 Changes in the income distribution and the effectiveness of monetary policy

This section investigates the impact of changes in the income distribution on the effectiveness of monetary policy. With this in mind, we simulate the response of aggregate prices to a monetary shock in two counterfactual calibrations of the model with different levels of income inequality. In the first counterfactual, we reduce inequality so that the share of aggregate income held by households at the top decile of the income distribu-

Table 2: Cumulative inflation, relative to aggregate

	Bottom 5%	Middle 5%	96-99 %	Top 1%
6 months	0.993	1.059	0.952	0.874
12 months	1.003	1.036	0.964	0.898
18 months	1.004	1.023	0.974	0.917
24 months	1.004	1.015	0.982	0.934
30 months	1.003	1.009	0.988	0.948
36 months	1.002	1.006	0.992	0.959

Notes: The table reports the impulse responses of the household-specific price indices P_t^h for households at the bottom, middle, and 5% of the income distribution, and for households at the top 1% of the income distribution, expressed relative to the impulse response of the aggregate price index, P_t.

tion is reduced by one half. This change in the share of income held by top households would roughly correspond to taking income inequality in the US back to 1980 levels.[8] In the second counterfactual, we increase share of income in the hands of the top decile in national income by 50 percent. In each counterfactual, we rescale the share of income of all the households below the top decile proportionally. Specifically, in the counterfactuals we set the shares to

$$s_c^h = \begin{cases} \alpha_c \times s_b^h & if\ h \in \text{top decile} \\ s_b^h \frac{1-s_c^{10\%}}{1-s_b^{10\%}} & else \end{cases},$$

where s_b^h and s_c^h are the baseline and counterfactual shares of aggregate spending by households of type h, respectively. We set $\alpha_{c_1} = 0.5$ in the first counterfactual and $\alpha_{c_2} = 1.5$ in the second counterfactual.

To conduct these counterfactuals, we calibrate the household-specific productivities A_c^h to match the desired income distribution, while leaving aggregate income unchanged. More generally, one could imagine that the shares would also change as we change each household's income. To reflect these counterfactual changes we proceed in two steps.

[8]Between 1980 and 2014, the share of US income held by the top 10% increased from 34% to 47%, and the share of income held by the top 1% increased from 10% to 20% according to the World Inequality Database (Alvaredo et al., 2016).

First, for each product category j, we split the households in the CES into consumption percentiles and perform a local polynomial regression of ω_j^h on s_b^h. Second, we calculate the predicted value of ω_j^h at the counterfactual shares s_c^h. We use these predicted values as the sectoral expenditure shares for households with counterfactual income A_c^h.

Figure 8 plots the impulse responses of the aggregate price index and of output in the two counterfactuals to a monetary shock that increases the nominal interest rate by 0.125 basis points on impact. The figure shows that prices are more responsive to monetary shocks in the baseline than in the counterfactual with more income inequality, and less responsive than in the counterfactual with less inequality. This is expected, since households at the top of the income distribution spend more of their income in sectors with more sticky prices. Prices decline by about 3.5% more in the counterfactual model with low income inequality for every horizon up to 24 months. However, the magnitude of the difference between the impulse responses of output is negligible. We conclude that realistic changes in inequality do not substantially alter how aggregate prices and output respond to monetary policy.

Figure 8: Response of the aggregate CPI and output to a monetary shock: Baseline vs. counterfactual income distributions

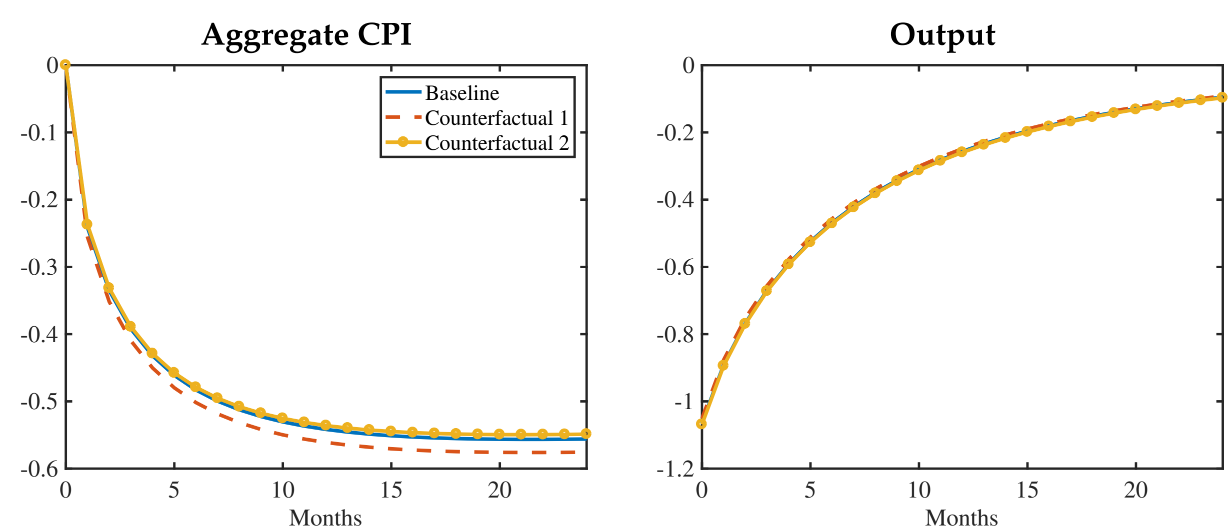

Notes: This figure plots the impulse responses of the aggregate price indices and output in the baseline calibration and in counterfactual a calibrations described in Section 5.2.3.

6 Conclusion

It has been known since at least Engel (1857, 1895) that households with different incomes consume different goods. This paper documents two novel patterns in how consumption baskets differ: in the United States, households at the top of the income distribution consume more sticky-priced goods and face substantially lower overall inflation volatility than households in the middle of the income distribution. Since the price stickiness, the volatility, and the response of prices to monetary policy differs across goods categories, these patterns suggest distributional consequences of monetary policy shocks. Because the prices of goods consumed by the high-income households are less responsive to monetary shocks, the overall CPIs of those households will react less to those shocks. We document both empirically and quantitatively that this is indeed the case. The estimated impulse responses to monetary shocks identified using the narrative approach of Romer and Romer (2004) show that CPIs of the high-income households react by about 38% less to a given monetary policy shock than CPIs of middle-income households 36 months after the shock. We then set up a multi-sector, heterogeneous-household model with sticky prices, parameterizing it to the observed sectoral heterogeneity in price stickiness and household heterogeneity in consumption baskets. In the model, the CPIs of high-income households respond 13% less to a monetary shock than the CPIs of middle-income households after 12 months.

References

Almås, Ingvild, "International Income Inequality: Measuring PPP Bias by Estimating Engel Curves for Food," *American Economic Review*, 2012, *102* (2), 1093–1117.

Alvaredo, Facundo, Anthony Atkinson, Lucas Chancel, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman, "Distributional National Accounts (DINA) Guidelines : Concepts and Methods used in WID.world," December 2016. WID.world WORKING PAPER SERIES N ° 2016/1.

Argente, David and Munseob Lee, "Cost of Living Inequality during the Great Recession," June 2015. Kilts Booth Marketing Series Paper 1-032.

Auclert, Adrien, "Monetary Policy and the Redistribution Channel," May 2017. mimeo, Stanford University.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra, "Regional Heterogeneity and Monetary Policy," August 2017. Forthcoming, *Quarterly Journal of Economics*.

Bernanke, Ben S., Jean Boivin, and Piotr Eliasz, "Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach," *Quarterly Journal of Economics*, 2005, *120* (1), 387–422.

Boivin, Jean, Marc P. Giannoni, and Ilian Mihov, "Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data," *American Economic Review*, March 2009, *99* (1), 350–84.

Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 1983, *12* (3), 383–398.

Christiano, Lawrence J., Mathias Trabandt, and Karl Walentin, "DSGE Models for Monetary Policy Analysis," in Benjamin M. Friedman and Michael Woodford, eds., *Handbook of Monetary Economics*, Vol. 3, Elsevier, 2010, chapter 7, pp. 285–367.

Clayton, Christopher, Xavier Jaravel, and Andreas Schaab, "Heterogeneous Price Rigidities and Monetary Policy," May 2018. Mimeo, Harvard and LSE.

Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong, "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications," *American Economic Review*, March 2015, *105* (3), 993–1029.

_ , _ , Lorenz Kueng, and John Silvia, "Innocent Bystanders? Monetary policy and inequality," *Journal of Monetary Economics*, June 2017, *88* (Supplement C), 70 – 89.

Cravino, Javier and Andrei A. Levchenko, "The Distributional Consequences of Large Devaluations," *American Economic Review*, November 2017, *107* (11), 3477–3509.

Doepke, Matthias and Martin Schneider, "Inflation and the Redistribution of Nominal Wealth," *Journal of Political Economy*, December 2006, *114* (6), 1069–1097.

Engel, Ernst, "Die Produktions- und Ernteerträge und der Getreidehandel im preussischen Staate," *Zeitschrift des Königlichen preussischen statischen Bureaus*, 1857, *2*, 249–89.

_ , "Das Lebenskosten belgischer Arbeiterfamilien früher und jetzt," *Bulletin de Institut International de Statistique*, 1895, *9*, 1–124.

Jaravel, Xavier, "The Unequal Gains from Product Innovations: Evidence from the US Retail Sector," April 2017. mimeo, LSE.

Kaplan, Greg and Sam Schulhofer-Wohl, "Inflation at the household level," *Journal of Monetary Economics*, November 2017, *91*, 19–38.

_, Benjamin Moll, and Giovanni L. Violante, "Monetary Policy According to HANK," *American Economic Review*, March 2018, *108* (3), 697–743.

Kim, Seongeun, "Quality, Price Stickiness, and Monetary Policy," April 2018. mimeo, Sejong University.

Nakamura, Emi and Jón Steinsson, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464.

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

Redding, Stephen J. and David E. Weinstein, "Measuring Aggregate Price Indexes with Demand Shocks: Theory and Evidence for CES Preferences," 2016. NBER Working Paper 22479.

Romer, Christina D. and David H. Romer, "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, September 2004, *94* (4), 1055–1084.

Williamson, Stephen, "Monetary policy and distribution," *Journal of Monetary Economics*, 2008, *55* (6), 1038–1053.

Wong, Arlene, "Transmission of Monetary Policy to Consumption and Population Aging," April 2016. Mimeo, Princeton University.

ONLINE APPENDIX
(NOT FOR PUBLICATION)

Appendix A Data Appendix

A.1 Constructing percentile-level expenditure weights

A.1.1 Consumer Expenditure Survey

We use data from the Consumer Expenditure Survey (CES) to obtain the expenditure weights of consumers. The CES data are collected by the Census Bureau, and cover expenditures, income, and demographic characteristics of households in the United States. The CES is the primary source of data for constructing the weights for the US Consumer Price Index.

The CES contains two modules, the Diary and the Interview. The Diary is designed to measure expenditures on daily items, such as groceries, personal products, and other frequent purchases. The Interview is designed to measure large or durable expenditures, such as major appliances, vehicles, and other large infrequent purchases. The Diary records household spending for two consecutive survey reference weeks, while the Interview records purchases over the previous three months.

For each survey, we make use of expenditure, income, and characteristics files in computing expenditure weights. In the expenditure files, the CES collects household expenditures on about 600 Universal Classification Code (UCC) categories. Questions such as "How much did you spend on babysitting in the last quarter" are asked in the survey and the corresponding responses are saved in *UCC 340210 babysitting and child care*. Overall, there are questions on about 350 UCC categories in the Interview module, and on 250 UCCs in the Diary module. Income files record detailed information on household monthly income from different sources, such as wages and salaries, or interest and dividends. Characteristics files record demographic characteristics data for each member of the household, such as education, gender, race, etc. Income variables, which contain annual values for the 12 months prior to the interview month, are also included in the characteristics files.

Diary and Interview modules survey different households each year, so a household in the Diary will not appear in the Interview and vice versa. Thus we could never observe the full consumption profiles of an actual household and we could not compute expenditure shares for an actual household. Rather, we aggregate households into percentiles and work with the percentile-level household expenditure shares.

A.1.2 Constructing the concordance

The in-scope expenditures for CPI could be divided into 8 major groups, 70 expenditure classes, 211 item strata (item level) and 303 entry level items (ELI). CPI uses the item strata -- e.g. *SEFT04 Spices, seasonings, condiments, sauces* -- as the elementary level of its expenditure weights and price index calculation. Within each item stratum, one or more substrata are defined as ELIs, which are the ultimate sample units for products. For example, there are four ELIs under item *SEFT04*: *FT041 Salt and other seasonings and spices*, *FT042 Olives, pickles and relishes*, *FT043 Sauces and gravies* and *FT044 Other condiments*.

Using CES data to compute the item-level and ELI-level expenditure weights from

CES, we need a concordance between the UCC categories, item strata codes and the ELIs. The concordance is constructed by following the BLS document *"CPI requirements for CE"* Appendix B. The CES collects household expenditures on about 600 Universal Classification Code (UCC) categories, which could be concorded to 303 ELIs. To combine the expenditure weights with the frequency of price adjustment data from Nakamura and Steinsson (2008), we look at a subsample of 265 ELIs. And we could further aggregate the 265 ELIs to 178 item strata.

A.1.3 Compiling the expenditure, income, and characteristics files

To obtain the expenditure shares at the detailed product category level for households at different percentiles of the income distribution, we take the following steps.

In the first step, we put together the quarterly expenditure, income, and characteristics files from the Interview survey. With the compiled interview data, for each household, we could observe its interviewed month and year, monthly expenditures on the UCC categories in the previous three months as well as annual income for the 12 months prior to the interview. One thing to note is that respondents are asked to report expenditures made since the first of the three months prior to the interview month. For example, if a household is interviewed in February of 2015, they are reporting expenditures for November and December of 2014, and January of 2015. Thus, to produce a 2014 annual estimate based on expenditures made in 2014 (calendar period), one needs to access five collection-quarter files, the first quarter of 2014 through the first quarter of 2015.

By the same token, we put together the expenditure, income, and characteristics files from the Diary survey. For each household in the Diary survey, we are able to observe its weekly expenditure on the detailed UCC categories and its annual income for the 12 months prior to the interview. Then we append the compiled Interview data file to the compiled Diary one to get the whole sample of UCCs.

A.1.4 Adjusting the expenditure values

In the second step, we make several adjustments to the collected expenditures in order to meet the BLS's requirements for the creation of CPI expenditure weights. The adjustments are made following the BLS document *"CPI Requirements of CE"*.

Housing

Two adjustments are made to housing categories.

- **Owners' equivalent rent of primary residence**
  UCC categories only collect the value of the house, its property taxes, real estate fees, and mortgage interests. Houses and other residential structures are capital goods and should not be considered as CPI items. Interest costs (such as mortgage interest), property taxes and most maintenance costs, are part of the cost of the capital good and are not consumption expenditures either. All of these are not useful in computing the expenditure weights for the item *Owners' equivalent rent of primary residence*.

According to the BLS document *"How the CPI measures price change of Owners' equivalent rent of primary residence (OER) and Rent of primary residence (Rent)"*, the expenditure weight in the CPI market basket for *Owners' equivalent rent of primary residence (OER)* is based on the following question that the CES asks consumers who own their primary residence:
"If someone were to rent your home today, how much do you think it would rent for monthly, unfurnished and without utilities?"
CES collects the household responses to this question and saves them in the variable *RENTEQVX* in characteristics files. We construct an artificial UCC code "999999" to store the values of variable *RENTEQVX*, which provides the household expenditure on the owners' equivalent rent of primary residence.

- **Homeowner insurance/maintenance/major appliance**
  The BLS adjusts the expenditures on homeowner insurance, maintenance, and major appliances to separate the consumption components of those expenditures from the investment component. The BLS uses a factor of 0.43 to account for the consumption portion of a homeowner's total expenditure on these housing categories. The factor is based on the likelihood that renters will purchase these types of appliances and perform these types of home maintenance and improvement. Thus, to reflect the consumption portion of a homeowner's total expenditure on housing insurance, maintenance, and major appliances, we multiply the expenditures on the corresponding UCC categories by 0.43.

Medical care

The BLS uses the National Health Expenditure (NHE) tables produced by the Center for Medicare and Medicaid Services (CMS) to calculate the factors that redistribute the weights from private health insurance and Medicare premium to medical care services. Unfortunately, we do not have access to the underlying formulas the BLS used to calculate these factors. By way of approximation, we take the redistributing factors from the NHE *Table 20 Private Health Insurance Benefits and Net Cost; Levels, Annual Percent Change and Percent Distribution, Selected Calendar Years 1960-2015*.[9]

We redistribute the expenditures from private health insurance and Medicare premiums related UCC categories to health care services categories, such as nursing homes and adult day services, by using factors obtained from the table mentioned above. Note that medical reimbursements are allocated across all households to smooth the household expenditures on medical expenses. That is to say, a household may be reimbursed even during a period in which they had no medical expenses.

Transportation

- **Used cars**
  Expenditures on used cars and trucks should only reflect dealer value added. There-

[9] For more details see the link https://www.cms.gov/research-statistics-data-and-systems/statistics-trends-and-reports/nationalhealthexpenddata/nhe-fact-sheet.html

fore, the expenditure weights on used cars and trucks should be determined by spending on used cars and trucks, minus trade-in value of vehicles and other sales of consumer-owned vehicles.
CES does not provide data on trade-in values of vehicles (*UCC 450116* and *450216*) and other sales of consumer owned vehicles (*UCC 860100* and *860200*). Thus, we take the expenditure weight on used cars and trucks from the BLS released table *Relative Importance of Components in the Consumer Price Index* to recover the ratio of trade-in values and other sales of vehicles to spending on used cars and trucks, and we find the ratio is around 1/2. Thus, we reduce the spending on used cars and trucks to half to reflect only the dealer value added.

- **Gasoline**
Gasoline expenditures are not allocated into categories (regular, premium, midgrade, etc.) at collection. To distribute the total gasoline expenditures (*UCC 470111*) amongst the gasoline ELIs (*TB011 Regular Unleaded Gasoline,TB012 Midgrade Unleased Gasoline and TB013 Premium Unleased Gasoline*), the BLS constructed the distribution factors from expenditure habits in each primary sampling unit (PSU).
However, we don't have access to the expenditure habits of each PSU. Instead, we follow Nakamura and Steinsson (2008), and allocate the expenditures on gasoline to regular, premium, midgrade categories equally.

A.1.5 Aggregating households into percentiles

In the third step, we aggregate households into percentiles. Because the Interview and the Diary survey different households, we sort the households into percentiles in two sub-steps. First, we aggregate the households in the Interview survey into percentiles based on imputed household annual income before tax, and then find the income cut-offs for each percentile. Second, we use the Interview survey income cut-offs to divide households from the Diary survey into percentiles. In this case, each household in our data sample has been sorted into a percentile. We could get similar results by using income cutoffs from the Diary survey to aggregate households in the Interview survey into percentiles.

The CES data start to include the imputed income since 2004. Before that it only publishes income data collected from households that are complete income reporters. Households are defined as complete reporters if they report one of the major sources of income, such as wages and salaries, Social Security income, or self-employment income. However, even a complete reporter might not provide information on all sources of income they indicate they received. Thus, in cases when the values of income are not reported, imputation allows them to be estimated. We sort households into percentiles based on the imputed household income before tax, which is only available since 2004. Because of this, therefore, we could only compute the percentile-level expenditure weights since 2004.

Table A1 reports the income cutoffs and average incomes in the selected quantiles of the income distribution.

Table A1: Income cutoffs and averages for selected quantiles of the income distribution in the CES

	Cut-offs			
	Lower	Upper	Median	Mean
Bottom 5%	-23,297	5,838	2,343	2,450
Middle 40-60%	36,504	62,808	48,828	48,969
96-99%	212,148	332,196	249,677	253,900
Top 1%	332,279	846,706	392,148	414,011

Notes: The table the range and the averages of the incomes in selected quantiles of the income distribution in the CES data.

A.1.6 Calculating the expenditure shares

In the final step, we calculate the expenditure shares at the detailed product category level for households at different income percentiles.

First, we calculate the average expenditure for each detailed UCC category for households at different income percentiles. Note that there is a distinction between survey period and expenditure reference period in the interview survey, as the CES collects household spending in the three months prior to the interview month. This distinction will affect the estimation procedure for producing household average expenditure during a calendar year. For example, households interviewed in February will report their spending for November and December of 2014 and January of 2015. Thus, to compute the average value for expenditures made on a certain UCC category during year 2015, they only contribute one month (January) of the expenditures they made during the expenditure reference period to the calculation. While households interviewed in May report their expense for February, March and April of 2015 and could contribute all their expenditures to compute the average expenditure this household made during 2015. To reflect the number of months a household can contribute to the mean value of a calendar year, we follow the BLS to create a variable called MO_SCOPE. In the above example, MO_SCOPE=1 for households interviewed in February and MO_SCOPE=3 for households interviewed in May. There is no such distinction between the survey period and expenditure reference period in the Diary. We multiply each weekly expenditure by 13 to get a corresponding quarterly expenditure. As there is no lag between the survey period and the expenditure reference period, the number of months households in the Diary survey contribute to estimate of the mean value is 3, i.e. MO_SCOPE=3. We could also interpret MO_SCOPE as the number of months a household reports expenditures during a calendar year.

Following the BLS manual, we use the formula below to calculate the average expenditure for each UCC category k at each percentile h. First, for household i at percentile h, we sum over all the spending it made on good k during the calendar year. Second, we weight total expenditures made by household i in percentile h on good k up by its household-specific sampling weight. Third, we sum up the weighted household expenditures on good k over all the households at percentile h. Fourth, we divide the sum of weighted household expenditures on good k at percentile h by the sum of the weighted number of months household at percentile h reported expenditures during the calendar year, to get the monthly average income on good k of household at percentile h. Then multiplying the monthly average expenditure by 12, we get the annualized average expenditure for each UCC category k at percentile h:

$$\bar{X}_k^h = \frac{\sum_i FINLWT_i^h \cdot \sum_t C_{i,k,t}^h}{\sum_i FINLWT_i^h \cdot MO\_SCOPE_i^h} \times 12$$

where $FINLWT_i^h$ is the sampling weight for household i at income percentile h, $C_{i,k,t}^h$ is the expenditure on good k of household i at income percentile h during month t, and $MO\_SCOPE_i^h$ denotes number of months household i reports expenditures during a calendar year.

Second, we take percentile-level average expenditure for each UCC from above, and then aggregate according to the constructed concordance between UCC categories and ELIs (or Item Strata) to get percentile-level household average expenditure $\bar{X}_j^h$ for each 265 ELIs (or 178 items) and the corresponding percentile-level expenditure share $\omega_j^h = \frac{\bar{X}_j^h}{\sum_j \bar{X}_j^h}$.

A.2 Constructing income-percentile-specific CPIs

A.2.1 Item-level consumer price data

To construct the income-percentile-specific consumer price indices (CPI), we need to combine the percentile-level expenditure share data computed above with the micro consumer price data. We obtain the consumer price data from the BLS. Each month, the BLS releases the consumer price index at all levels of aggregation. Each price index has a unique identifier called series id, CUUR0400AA0 for example. The series id can be broken down to: CU–survey abbreviation–current series, U–season code–seasonal unadjusted, R–periodicity code–monthly, 0400–area code–Western Urban and AA0–item code–all items. We use the U.S. city average, all urban consumers, seasonally adjusted item-level monthly price indices to construct the monthly income-percentile-specific CPIs.

A.2.2 Concordance between old and new series

The revised consumer price data were introduced by the BLS in 1998, and the revision included an updated and revised item structure. For example, there were only 7 major groups of goods and services before 1997 and in 1998, a new group *Education and Communication* was created and the new group included components previously included in the *Recreation* and *Housing* groups. Here, we refer to the revised item structure as the new series and to the pre-revised item structure as the old series. Micro consumer price data are provided in the old series before 1997, and in new series since 1997.

To combine the item-level consumer price data from the old series with the expenditure share data, we manually construct a concordance from the new series to the old series at the item level. Note that there are some new series items that are more aggregated than the old ones, and in these cases one item in the new series is concorded to multiple items in old series. To deal with it, we construct a concordance weight by using the expenditure weight taken from the BLS table *Relative Importance of Components in the Consumer Price Index*. One example is as follows. Item *SEFF01 Chicken* in the new series is concorded to *SE0601 Fresh whole chicken* and *SE0602 Fresh/Frozen chicken parts* in the old series. We find that the average expenditure during years 1987 to 1989 on the two items are 0.152% and 0.220% respectively, and thus we assign the concordance weights based on their relative expenditure weights on the two items. The 265 new series items are concorded to 165 old series ones.

item code (new)	item name (new)	item code (old)	item name (old)	exp weight	concordance weight
SEFF01	CHICKEN	SE0601	FRESH WHOLE CHICKEN	0.152	0.409
SEFF01	CHICKEN	SE0602	FRESH/FROZEN CHICKEN PARTS	0.220	0.591

A.2.3 Aggregation formula

We follow the BLS manual *"Chapter 17. The Consumer Price Index"* in constructing the income-percentile-specific CPI. The formula can be written as follows:

$$PIX_t^h = PIX_v^h \cdot \sum_{j \in J} (\omega_{j,\beta}^h \times \frac{P_{j,t}}{P_{j,v}}),$$

where:
$PIX_t^h =$ consumer price index for household at percentile h at time t
v = pivot year and month, usually December, prior to the month when expenditure weights from reference period (β) are first used in the CPI
β = predetermined expenditure reference period
$P_{j,t}$ = price of item j at time t
$\omega_{j,\beta}^h$ = expenditure weights of household at percentile h on item j during the predetermined expenditure reference period β .

The BLS periodically updates its expenditure weight reference period. Historically, it updated approximately every ten years, and since 2002, it adopted a biennial rotation schedule to update the expenditure weight reference period. We follow the BLS expenditure reference period schedule after 2004, and prior to that, we use the 2004 percentile-level expenditure weights to construct the income-percentile-specific CPI. As mentioned in A.1.5, this is due to the availability of the imputed household income before tax. We have computed the pre-2004 aggregate CPI by taking the official expenditure weights from BLS table *Relative Importance of Components in the Consumer Price Index* for the pre-2004 expenditure reference period. And comparing it with the aggregate CPI constructed by using 2004 aggregate weights, we find the two CPI series are almost identical.

Due to the revision of item structure in 1998, we have to construct the income-percentile-specific CPI separately in two periods. We use old series item-level micro price data to compute the income-percentile-specific CPI for the period 1969m1-1997m12 and new series price data for the post-1998 period. In the year 1997, the BLS released item-level micro prices in both old and new series, which allows us to bridge the two periods by using one of the months in 1997 as the pivot month (based period) for the second period. We used both the old and new series micro price data to construct the aggregate CPIs in 1997 and found that they give us similar results in (log) price terms. We choose 1997m12 as the first pivot month for the construction of the post-1998 income-percentile-specific CPIs.

Table A2: Reference periods

pivot month (v)	reference period (β)	PIX (t)
1969M1	2004	1969-1997(Old series)
1997M12	2004	1998-2005(New series)
2005M12	2004-2005	2006-2007(New series)
2007M12	2006-2009	2008-2009(New series)
$\vdots$	$\vdots$	$\vdots$
2015M12	2012-2015	2016-2017(New series)

Notes: This table lists the reference periods used to construct the CPI.

A.3 Categories with the largest expenditure share differences

Table A3 reports the 10 categories with the largest differences in expenditure shares between the top 1% and the middle 20% of households.

Table A3: Expenditure share differences, frequency of price adjustment, and volatility of price changes

	Income percentile			Regular	Price	St.
Category	40-60	100	Difference	Price Change	Change	Dev.
Top 10, larger expenditure shares by middle class						
Gasoline (all types)	0.084	0.038	-0.046	87.71	87.74	0.208
Electricity	0.050	0.025	-0.025	38.14	38.14	0.035
Limited service meals and snacks	0.037	0.018	-0.018	6.13	7.00	0.009
Wireless telephone services	0.032	0.014	-0.018	13.00	13.00	0.044
Motor vehicle insurance	0.039	0.021	-0.018	8.16	8.16	0.025
Hospital services	0.037	0.024	-0.013	6.26	6.26	0.014
Cable and satellite television and radio service	0.024	0.012	-0.013	12.35	12.83	0.015
Used cars and trucks	0.028	0.016	-0.012	100.00	100.00	0.052
Prescription drugs	0.022	0.011	-0.011	15.03	15.09	0.015
Cigarettes	0.012	0.001	-0.010	23.17	33.59	0.073
Mean				31.00	32.18	0.049
Median				14.02	14.04	0.030
Top 10, larger expenditure shares by top 1%						
College tuition and fees	0.012	0.051	0.039	5.77	5.77	0.018
Child care and nursery school	0.006	0.030	0.024	6.91	6.91	0.011
Elementary and high school tuition and fees	0.002	0.025	0.023	6.23	6.23	0.013
Watches	0.001	0.021	0.021	3.06	19.83	0.028
Airline fare	0.008	0.028	0.020	59.84	59.84	0.062
Domestic services	0.002	0.019	0.017	4.31	4.31	0.014
Club dues and fees for participant sports and group exercises	0.006	0.022	0.016	8.57	12.56	0.017
Other lodging away from home including hotels and motels	0.007	0.023	0.015	41.73	42.75	0.034
New vehicles	0.048	0.057	0.009	18.89	19.45	0.014
Admissions	0.005	0.013	0.008	8.07	8.39	0.017
Mean				16.34	18.60	0.023
Median				7.49	10.47	0.017

Note: This table reports the product categories with the largest differences in expenditure shares between the middle (40th-60th percentiles) and the top 1% of the income distribution, the frequency of price changes, and the standard deviation of 12-month log price changes for those products.

Appendix B FAVAR evidence

This appendix presents an alternative method to estimate the impulse responses of income-specific CPIs to monetary policy shocks: the Factor-Augmented Vector Autoregression (FAVAR) approach of Bernanke et al. (2005) and Boivin et al. (2009). Let there be a large number of economic series, whose behavior is driven by a vector of common components. This vector includes monetary policy in the form of the Federal Funds rate i_t, and a small number of unobserved common factors $\mathbf{F}_t$. The joint evolution of the Federal Funds rate and the vector of factors, $\mathbf{C}_t$, is characterized by a VAR:

$$\mathbf{C}_t \equiv \begin{bmatrix} \mathbf{F}_t \\ i_t \end{bmatrix},$$

$$\mathbf{C}_t = \Phi(L)\mathbf{C}_{t-1} + \mathbf{v}_t, \tag{B.1}$$

where $\Phi(L)$ is a lag polynomial, and $\mathbf{v}_t$ is an i.i.d. error term.

The vector $\mathbf{F}_t$ is unobservable. What is observed is a large number of economic series $\mathbf{X}_t$. The FAVAR approach assumes that this set of economic series is characterized by a factor model:

$$\mathbf{X}_t = \mathbf{\Lambda}\mathbf{C}_t + \mathbf{e}_t, \tag{B.2}$$

where Λ is the matrix of factor loadings. This representation provides a great deal of parsimony because in practice $\mathbf{X}_t$ includes hundreds of series, whereas the dimensionality of the vector of common factors $\mathbf{F}_t$ is typically small: in the Boivin et al. (2009) implementation there are 5 common unobserved factors. The significant benefit of estimating model (B.1)-(B.2) is that it yields impulse responses of each of the hundreds of series contained in $\mathbf{X}_t$ to shocks to the elements of $\mathbf{C}_t$, including monetary policy.

In our application of this approach, the vector $\mathbf{X}_t$ includes the 100 income-percentile-specific consumption price indices, as well as the additional variables included by Bernanke et al. (2005) and Boivin et al. (2009), such as sector-level industrial production, employment and earnings, and industry-product-level PPI series. The time frequency is monthly, and the time period is 1978m1-2008m12. Boivin et al. (2009) present a detailed evaluation of the performance of the FAVAR model. Here, we focus on the element new in our paper, namely the impulse responses of income-specific CPIs to monetary policy shocks.

The FAVAR produces 100 of those impulse responses, one for each income percentile. Figure A.1 plots those impulse responses for selected percentiles. The monetary policy shock is a 25-basis-point increase in the Federal Funds rate on impact, thus a contraction. The consumption price indices of the high-income households react substantially less to monetary policy shocks than those for the middle of the income distribution. The difference is economically meaningful. After 12 months, the top-1% households' CPI responds by 34% less, and the 96-99th percentile households by 22% less, than the CPI of the households in the middle of the income distribution (40-60th percentiles). After 24 months, the differences are still 12% and 6%, respectively.

A well-known feature of the VAR impulse responses of prices to monetary shocks is that the confidence intervals are wide, and it is often not possible to reject a zero impact of a monetary shock on aggregate CPI. This is the case in the Boivin et al. (2009) FAVAR

Figure A.1: Income-specific CPI impulse responses to a monetary policy shock

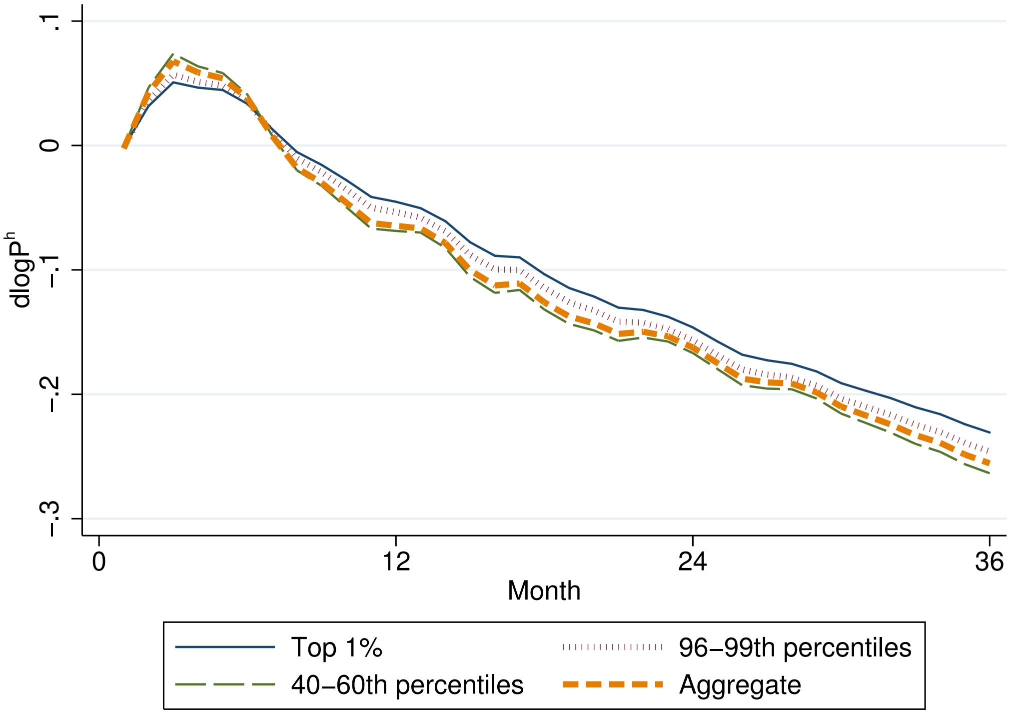

Notes: This figure plots the impulse responses of income-specific price indices to a monetary policy shock, estimated using a FAVAR.

model that forms our baseline analysis. However, our main object of interest is not the overall response of prices to a monetary shock, but rather the differential response of the CPIs of different households. Figure A.2 plots the difference in the impulse responses between the CPI of the top 1% and the CPI of the middle 20% of the income distribution (left panel), and the difference between the top 1% and aggregate CPI (right panel). Both panels include the 90% bootstrapped confidence intervals. The difference between impulse responses is significant at the 10% level for most of the lags between 8 and 21 months.[10]

[10]Note that the impulse is a monetary contractions, and thus the changes in the CPIs are negative after an initial few months. Since the top-income CPIs respond by less in absolute terms, the difference between the top- and middle-income CPIs is positive.

Figure A.2: Differences in inflation changes between income groups

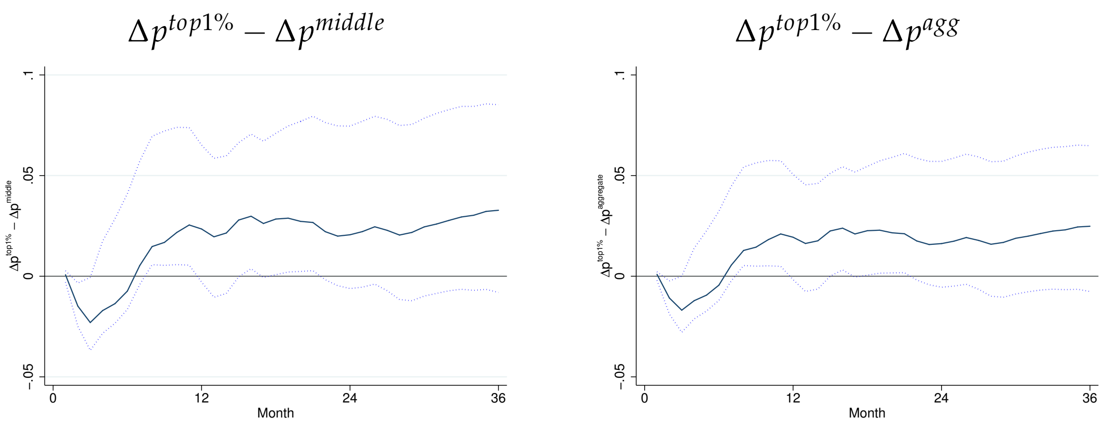

Notes: The left panel plots the difference between the impulse responses of the price index of the top 1% of households and the middle 20% of households to a monetary shock, while the right panel plots the difference between the impulse responses of the price index of the top 1% of households and the aggregate price index, along the 90% bootstrapped confidence intervals.

Appendix C Substitution bias

The results in this paper build on the assumption that changes in expenditure shares only have second order effects on inflation. Indeed, the Laspeyres index can be thought of as a first-order approximation to the change in the ideal price index, and thus we rely on the first-order approximation being suitable in this setting. To evaluate this assumption, this Appendix uses year-specific aggregate expenditure shares for each consumption Item from the CES data to construct Laspeyres and Paasche price indices. Since the ideal price index is in-between the Laspeyres and the Paasche, the difference between these two indices provides the upper bound on the bias induced by the first-order approximation.

Figure A.3 below summarizes these differences. It plots 12-month inflation rates for the aggregate CPI computed with Laspeyres and Paasche formulas. From 2004 onwards, we can obtain year-specific aggregate expenditure shares for each consumption Item from the CES data. The CES is the source of expenditure shares data used in the paper. Unfortunately, the product and income definitions in the CES are hard to harmonize prior to 2004. The right panel of Figure A.3 complements the CES data using year-specific aggregate expenditure shares from the BLS between 1987 and 2004 (these expenditure shares are only available from 1987). Both the CES- and the BLS-based measures show little difference between the Laspeyres and the Paasche inflation rates, which confirms that the substitution bias is indeed small in these years. Table A4 shows that the mean and the standard deviation of the Laspeyres and the Paasche inflation rates are an order of magnitude larger than the mean and standard deviation of the difference between the two measures. The correlation between the Laspeyres and the Paasche inflation rates is 0.99.

Note that the differences between the Laspeyres and the Paasche price indices likely overstate the importance of the substitution bias. Measured expenditure shares may change for reasons other than price changes, such as changes in the composition of households in the expenditure surveys, or changes in tastes across years (see Redding and Weinstein, 2016). In fact, as shown in Table A4, the standard deviation of the Paasche inflation is larger than of the Laspeyres inflation. Also, there are a number of occasions in which the Paasche inflation is larger than the Laspeyres inflation. This should not be the case if yearly changes in expenditure weights are solely due to substitution towards lower-inflation items.

Table A4: Comparison between the Paasche and Laspeyres price index inflation

	π^L		π^P		$\pi^P - \pi^L$		abs$(\pi^P - \pi^L)$		$Correl(\pi^P, \pi^L)$
	mean	SD	mean	SD	mean	SD	mean	SD	
1988-2004	2.95%	1.34%	3.11%	1.44%	0.16%	0.17%	0.17%	0.15%	0.99
2004-2016	2.08%	2.13%	2.06%	2.26%	-0.02%	0.40%	0.25%	0.32%	0.98

Notes: This table reports the mean and the standard deviation for the Laspeyres price index (π^L), Paasche price index (π^P), the difference between the two ($\pi^P - \pi^L$), and the absolute difference between the two (abs$(\pi^P - \pi^L)$). The last column reports the correlation between the Laspeyres and Paasche inflation rates. The inflation rates are defined as 12-month log changes in the price indices.

Figure A.4 uses year-specific expenditure weights by income level computed from the

Figure A.3: Aggregate Laspeyres and Paasche CPI inflation

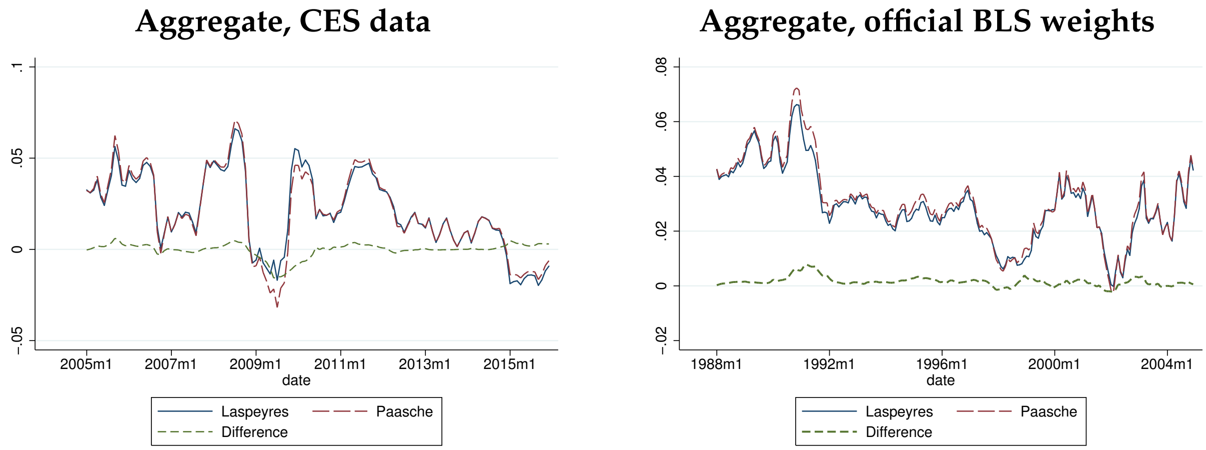

Notes: This figure plots the Laspeyres and Paasche indices, and the difference between the two, for aggregate 12-month CPI inflation. The left panel uses annual aggregate expenditure weights from the CES, the right panel from the BLS.

CES to construct Laspeyres and Paasche inflation for the households at the middle and at the top 1% of the income distribution. For the reasons mentioned above, we can only construct these series starting in 2004. Summary statistics for these measures are reported in Table A5. For each income group, the figure shows that the difference between the Paasche and the Laspeyres inflation is small compared to the overall inflation rates for both groups of households.

Table A5: Comparison between the Paasche and Laspeyres price index inflation, top- and middle-income households

Income	π^L		π^P		$\pi^P - \pi^L$		abs($\pi^P - \pi^L$)		$Correl(\pi^P, \pi^L)$
	mean	SD	mean	SD	mean	SD	mean	SD	
Top	1.88%	1.44%	1.94%	1.60%	0.05%	0.35%	0.27%	0.22%	0.98
Middle	2.16%	2.36%	2.12%	2.54%	-0.04%	0.50%	0.32%	0.38%	0.98

Notes: This table reports the mean and the standard deviation for the Laspeyres price index (π^L), Paasche price index (π^P), the difference between the two ($\pi^P - \pi^L$), and the absolute difference between the two (abs($\pi^P - \pi^L$)). The last column reports the correlation between the Laspeyres and Paasche inflation rates. The inflation rates are defined as 12-month log changes in the price indices.

Figure A.4: Laspeyres and Paasche CPI inflation by income level

Notes: This figure plots the Laspeyres and Paasche indices, and the difference between the two, for the middle 20% of the households (left panel), and the top 1% of the households (right panel) in the CES.