BANK OF CANADA
BANQUE DU CANADA

Staff Working Paper/Document de travail du personnel—2024-45

Last updated: November 21, 2024

Preferences, Monetary Policy and Household Inflation

by Geoffrey Dunbar

International Economic Analysis Department
Bank of Canada
gdunbar@bankofcanada.ca

DOI: https://doi.org/10.34989/swp-2024-45 | **ISSN** 1701-9397

Acknowledgements

The views in this paper are those of the author and no responsibility for them should be attributed to the Bank of Canada or its employees. This paper is loosely related to previous, uncirculated, work with my former research assistant, Shane Wood. In general, I thank Shane for his valuable assistance with cleaning the IRI data. I would like to thank IRI for making the data available. All estimates and analysis in this paper, based on data provided by IRI, are by the author and not by IRI.

Abstract

Household inflation can be decomposed into three terms that reflect nominal expenditure, real quantities and household preferences, using the money pump proposed by Echenique, Lee and Shum (2011). I quantify the importance of changes in household preferences on household inflation rates using 11 years of scanner data for 11,000 US households. I then analyze the effect of monetary policy on household inflation using the monetary policy shocks from Nakamura and Steinsson (2018). I find that monetary policy news shocks affect household inflation through the expenditure and preferences channels for 12 months from the date of the shocks, and that federal funds rate shocks affect inflation through the same channels at a horizon of 13–24 months. The results suggest that changes in household preferences are an important driver of inflation dynamics at the household level.

Topics: Inflation and prices; Monetary policy transmission
JEL codes: D12, E52, E58

Résumé

L'inflation subie par les ménages peut être décomposée en trois termes – la dépense nominale, les quantités réelles et les préférences des ménages – à l'aide d'une unité de mesure (money pump cost) proposée par Echenique, Lee et Shum (2011). Je quantifie l'incidence de l'évolution des préférences des ménages sur les taux d'inflation des ménages en me fondant sur 11 ans de données recueillies à l'aide de lecteurs de codes barres pour 11 000 ménages américains. J'analyse ensuite l'effet de la politique monétaire sur l'inflation subie par les ménages au moyen des chocs de politique monétaire définis par Nakamura et Steinsson (2018). Je constate que les chocs liés aux annonces de politique monétaire se répercutent sur l'inflation subie par les ménages à travers leurs dépenses et leurs préférences pendant 12 mois à compter du moment où ils se produisent. Quant aux chocs liés aux annonces de taux des fonds fédéraux, ils prennent 13 à 24 mois pour se répercuter sur l'inflation à travers les mêmes canaux. Les résultats portent à croire que les changements de préférences des ménages jouent un rôle important dans la dynamique de l'inflation du côté des ménages.

Sujets : Inflation et prix ; Transmission de la politique monétaire
Codes JEL : D12, E52, E58

1 Introduction

A central question in macroeconomics is how individuals or households react to monetary policy. Central banks in many, if not most, advanced economies have inflation targeting mandates, which suggests that their policy decisions influence changes in consumer prices. Whether this implies that consumer behaviour is affected by monetary policy is less clear. At the aggregate level, consumer price index (CPI) inflation is measured using expenditure weights that are updated infrequently, often only annually or biennially. An implication is that responses of CPI inflation to monetary policy decisions do not directly reflect consumer choice behaviour because fixed consumption baskets imply that consumer choices are held invariant. Using individual- or household-level fixed basket inflation data does not help because individuals and households rarely have fixed consumption baskets, even across adjacent months. Calculating inflation using a subset of goods repeatedly purchased is possible, however it is difficult to argue that inflation rates calculated this way reflect behavioural responses when the bundle of goods selected are precisely those for which there is no behavioural response.

In this paper, I show how household-level inflation rates reflect household preferences and that monetary policy affects inflation through its influence on households' consumption choices. I show that individual or household inflation can be decomposed into three separate terms—a nominal term, a real term, and a preference term—using the money pump cost proposed by Echenique, Lee, and Shum (2011). The money pump cost for adjacent time periods can be interpreted as a dollar (currency unit) measure of violations of the weak axiom of revealed preference (WARP).[1] While one can interpret violations of WARP as indicative of irrational preferences or evidence of measurement error, the interpretation favoured in this paper is that the money pump cost is the (implicit) cost for individuals or households for changing their preferences. The money pump decomposition also implies that an *as if* inflation measure can be calculated using all purchases in both reference periods and it explicitly quantifies how consumption choices affect inflation. I then examine how monetary policy affects each term of the household inflation decomposition and find, in particular, that monetary policy directly impacts the consumption choices made by households, which then has implications for household inflation.

Using scanner data for a panel of approximately 11,000 households from Nielsen IRI over the period 2001–2011, I calculate money pump costs for retail (mainly grocery) store expenditures for these households at a monthly frequency.[2] The IRI data collects Universal Product Code (UPC) level information for the household. The median monthly consumption for the product categories is approximately $50 per month. Like Echenique, Lee, and Shum (2011) and Smeulders, Cherchye, Spieksma, and Rock (2013), I find evidence of positive money pump costs for households in the sample—in an average month, roughly 30 to 40 percent of households have money pump costs greater than zero. Money pump costs are also somewhat persistent with an unconditional month-to-month auto-correlation of 0.35. On average, only 7 percent of items are purchased in the same quantity in adjacent periods, although this rises to 20 percent for identical items purchased in different quantities. This suggests that preferences may be a significant driver of household inflation. Indeed, I show that household inflation rates can differ by roughly 10 percentage points on average for households with positive money pump costs compared to households with negative money pump costs.

Finding heterogeneity in household inflation is consistent with recent evidence on household heterogeneity in inflation; see, for example, Kaplan and Schulhofer-Wohl (2017), Hobijn, Mayer, Stennis, and Topa (2009) and Hobijn and Lagakos (2005). One key difference between this literature and the current paper is the

[1] One interpretation of money pump costs is that it is the amount that could in principle be extracted from a consumer who violates the WARP.

[2] I use the Nielsen BehaviourScan scanner panel data comprising households in two US counties, Pittsfield, MA, and Eau Claire, WI, for 132 months obtained from IRI Marketing (Bronnenberg, Kruger, and Mela, 2008). As with any survey data, there are naturally questions regarding the representativeness of the sample. I report in Table 1 some demographic variables from the US Census and compare these with our sample counterparts for the two counties in our sample. This comparison does not affect the results we obtain for our sample but perhaps does bear on whether our results generalize, which is left as a consideration to the reader. We also note that one advantage of scanner data is that households are less likely to suffer from recall error.

role of household choice (behaviour) for household inflation heterogeneity.[3] Redding and Weinstein (2019) is a recent exception that does explicitly consider how consumer tastes affect price indices. However, they do not consider how economic factors influence consumer tastes, in particular how monetary policy affects household choice behaviour, as in this paper.

I examine how average household inflation responds to monetary policy decisions using identified monetary policy news and federal funds rate shocks from Nakamura and Steinsson (2018). Because it can be difficult to isolate the effect of monetary policy decisions from other coincident events, I use the three pass regression filter proposed by Kelly and Pruitt (2015) to leverage the cross section of households to construct factors for the shocks. The first stage of the three pass regression filter is time-series regressions for each household, which capture the time-series correlation between the household variables and the shocks. The second stage of the three pass filter uses cross-section regressions on the first stage coefficients to construct factors for the shocks for each period. The second stage regressions include a time fixed effect, which captures time variation orthogonal to the first stage coefficients for the shocks, i.e., orthogonal to the monetary policy shocks. The second stage factors are constructed separately for each term of the inflation decomposition, which isolates the channel through which the shocks affect overall inflation.

The results in this paper suggest that household inflation responds to the policy news shock through the expenditure and behaviour channels for a 12-month horizon. Household inflation also responds to the federal funds rate through the same channels at a horizon of 13–24 months from the time of the monetary policy decision. This suggests that policy statements by central banks affect household behaviour at a shorter horizon than policy rate changes do. The results also show the importance of the behaviour channel for the transmission of monetary policy. I find that the behaviour and expenditure channels tend to have offsetting effects, with the former channel having a typically larger effect. Ignoring the behaviour channel, I find evidence that monetary policy shocks tend to increase inflation through the expenditure channel. This suggests that the well-known price puzzle noted by Sims (1992) (and given its moniker by Eichenbaum, 1992) may be, at least partly, explained by the effect of monetary policy on household behaviour.

The dispersion of household inflation rates, measured by the inter-quartile range, is also affected by the interest rate factor after 12 months with positive (negative) federal funds rate shocks leading to an decrease (increase) in the inter-quartile range of inflation. Lauper and Mangiante (2021) calculate inflation rates for different subgroups of US households and find heterogeneous responses to monetary policy contractions across these household subgroups. In particular, they find that household heterogeneity can mute the impact of monetary policy for inflation. A similar conclusion is reached by Cravino, Lan, and Levchenko (2020), who find that price stickiness for consumption is positively related to household income. The results in this paper point to a similar importance for household heterogeneity. However, I also find some evidence that the dispersion of household inflation may decrease (increase) because of positive (negative) policy news shocks at a horizon of up to 12 months. This suggests that monetary policy communication may be consumed, or interpreted, differently across households and that the communications channel for central banks may be effective as a tool for managing household inflation inequality in the short run.

The remainder of the paper proceeds as follows. Section 2 reviews Echenique, Lee, and Shum (2011) and shows how the money pump cost proposed in that paper relates to household inflation. Section 3 describes the scanner data used in this paper and provides summary information for the sample composition of the households and the consumption items included. Section 4 provides unconditional time-series charts outlining the contribution of the three channels—nominal expenditure, real expenditure and money pump costs—for household inflation. Section 5 examines the role of monetary policy for average household inflation and the dispersion of household inflation. Section 6 concludes.

[3]In this paper, I use the term "choice" and "behaviour" interchangeably unless otherwise indicated.

2 Methodology

Calculating household-level inflation rates is complicated by the fact that households rarely have a fixed consumption basket, even over adjacent periods. One option is to calculate inflation rates for the subsets of goods that are consumed by the household in both periods (e.g., Kaplan and Schulhofer-Wohl (2017)). An obvious concern with this approach is that changes in consumption bundles are unlikely to be randomly distributed and thus it is not clear that an inflation rate on the subset of goods is indicative of the inflationary pressure experienced by the household. In this section, I show that the money pump cost proposed by Echenique, Lee, and Shum (2011) can be used to calculate household inflation rates when the basket of goods changes over time.

The money pump cost is a unit of account measurement of violations of the weak (and potentially strong) axiom of revealed preference. In what follows, $\mathbf{p}_t^i$ is the (row) vector of prices faced by the household at time t and $\mathbf{x}_t^i$ is the (column) vector of quantities purchased by household i at time t. I label the unit of account as money (sometimes specifically as dollars). Echenique, Lee, and Shum (2011) define the money pump cost for the household, M_t^i, as:

$$\mathsf{M}_t^i = \mathbf{p}_{t-1}^i(\mathbf{x}_{t-1}^i - \mathbf{x}_t^i) + \mathbf{p}_t^i(\mathbf{x}_t^i - \mathbf{x}_{t-1}^i). \tag{1}$$

M_t^i is a money-metric measurement, and $\mathsf{M}_t^i > 0$ implies household i's consumption choices do not satisfy WARP.[4] Echenique, Lee, and Shum (2011) also argue that the money pump cost can be interpreted as the amount of money that would be earned by a "devious arbitrageur" who exploits the violation of WARP by buying the reverse transactions and reselling them to the household. By implication, violations of WARP are a cost to the household and can be interpreted as an additional purchase by the household. This implies that money pump costs should be included in measurements of household inflation. The money pump cost can be rearranged to decompose household inflation, π_t^i, into three terms:

$$\pi_t^i = \frac{(\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\mathbf{x}_{t-1}^i}{\mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i} = \underbrace{\frac{\mathbf{p}_t^i\mathbf{x}_t^i}{\mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i}}_{\text{nominal expenditure change}} - \underbrace{\frac{\mathbf{p}_{t-1}^i\mathbf{x}_t^i}{\mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i}}_{\text{Laspeyres quantity change}} - \underbrace{\frac{\mathsf{M}_t^i}{\mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i}}_{\text{money pump cost}}. \tag{2}$$

This definition of household-specific inflation, π_t^i, follows the Lasperyes construction since it holds the consumption bundle fixed at the previous period realization. It shows that household-specific inflation rates can be decomposed into three parts: nominal expenditure growth, a quantity change, and money pump costs deflated by previous period expenditure.

While the money pump cost term in Equation (2) specifically accounts for the role of consumer choices on household inflation, it is not the case that consumption bundle differences necessarily impact household inflation. In the case where $\mathbf{x}_t^i = \mathbf{x}_{t-1}^i$ (consumption bundles are identical), then $\mathsf{M}_t^i = 0$ and

$$\pi_t^i = \frac{(\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\mathbf{x}_{t-1}^i}{\mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i} = \frac{(\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\mathbf{x}_t^i}{\mathbf{p}_{t-1}^i\mathbf{x}_t^i}.$$

However, $\mathsf{M}_t^i = 0$ does not necessarily imply consumption baskets are fixed. In fact, there are a continuum of possible consumption baskets $\mathbf{x}_t^i$ and $\mathbf{x}_{t-1}^i$ that conceivably yield $\mathsf{M}_t^i = 0$ for any price vectors $\mathbf{p}_t^i$ and $\mathbf{p}_{t-1}^i$. If $\mathsf{M}_t^i = 0$ then inflation is simply the difference between nominal expenditure growth and the quantity index growth. Whenever $\mathsf{M}_t^i \neq 0$ this statement is not true and preferences directly affect household inflation rates by Equation (2).

The effect of money pump costs is not specific to only CPI measures of inflation. An alternative measure of inflation, the personal consumption expenditures (PCE) deflator, is also affected. The US Bureau of

[4] Echenique, Lee, and Shum (2011) also show the money pump cost can be applied to the generalized axiom of revealed preference (GARP). It is possible that choices that do not satisfy WARP do satisfy GARP. The key difference is that WARP does not allow indifference between two or more consumption bundles at fixed prices. In empirical applications, violations of WARP are typically accompanied by violations of GARP; see Echenique, Lee, and Shum (2011) and Smeulders, Cherchye, Spieksma, and Rock (2013).

Economic Analysis (BEA) calculates the PCE price index using a Fisher-Ideal geometric mean of a Laspeyres and Paasche relative prices. It is also possible to define household-specific inflation using the Paasche definition:

$$\frac{(\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\mathbf{x}_t^i}{\mathbf{p}_t^i\mathbf{x}_t^i} = \frac{\mathbf{p}_t^i\mathbf{x}_{t-1}^i}{\mathbf{p}_t^i\mathbf{x}_t^i} - \frac{\mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i}{\mathbf{p}_t^i\mathbf{x}_t^i} + \frac{\mathsf{M}_t^i}{\mathbf{p}_t^i\mathbf{x}_t^i}.$$

The key difference if the Paasche definition is used is that positive (negative) money pump costs increase (decrease) household-specific inflation rates. One can show that the geometric average of the Laspeyres and Paasche relative prices includes a money pump effect, which implies that the PCE deflator is also affected by money pump costs if $\mathsf{M}_t^i \neq 0$.[5] For the remainder of this paper, I focus on the effect of preferences for a Laspeyres price index.

Because the argument in this paper is that changing preferences is costly, a simple sanity check is to show that these costs are reflected in the national accounts. Rearranging the money pump cost measure, Equation (1), and subtracting $\mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i$ from both sides yields:

$$\mathbf{p}_t^i\mathbf{x}_t^i - \mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i = \mathsf{M}_t^i + (\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\mathbf{x}_{t-1}^i - \mathbf{p}_{t-1}^i(\mathbf{x}_{t-1}^i - \mathbf{x}_t^i), \tag{3}$$

which shows that nominal expenditure changes for a household (given by the left-hand side) also depend on M_t^i. I define, for period j, the population consumption $\mathbf{X}_j = \sum_{i=1}^{N_j} \mathbf{x}_j^i$, where N_j is the total population, the population money pump costs $\mathrm{M}_j = \sum_{i=1}^{N_j} \mathsf{M}_j^i$ and, for ease of exposition, I assume prices are common across households. Normalizing by the previous periods aggregate expenditure, $\mathbf{p}_{t-1}\mathbf{X}_{t-1}$, gives:

$$\underbrace{\frac{\mathbf{p}_t\mathbf{X}_t}{\mathbf{p}_{t-1}\mathbf{X}_{t-1}} - 1}_{\text{nominal PCE growth rate}} = \underbrace{\frac{\mathbf{X}_{t-1}(\mathbf{p}_t - \mathbf{p}_{t-1})}{\mathbf{p}_{t-1}\mathbf{X}_{t-1}}}_{\text{PCE price deflator}} + \underbrace{\frac{\mathbf{p}_{t-1}(\mathbf{X}_t - \mathbf{X}_{t-1})}{\mathbf{p}_{t-1}\mathbf{X}_{t-1}}}_{\text{PCE real growth}} + \underbrace{\frac{\mathrm{M}_t}{\mathbf{p}_{t-1}\mathbf{X}_{t-1}}}_{\text{money pump adjustment}}, \tag{4}$$

demonstrating that aggregate money pump costs are positively related to aggregate nominal personal consumption expenditures in national accounts data. While this is perhaps an interesting avenue of exploration, it is left for future research as the focus in the current paper is on the role of money pump costs for inflation. In addition, the scanner data used in this paper covers only a small subset of consumer purchases, and I do not have access to the microdata necessary to investigate this channel for a more representative sample of household expenditure.

While measuring the impact of household behavioural changes for inflation would seem a useful contribution, Equation (2) is also suggestive of a normative contribution for inflation-targeting central banks. Violations of WARP, $\mathsf{M}_t^i > 0$, lower household inflation, while households whose preferences satisfy WARP, $\mathsf{M}_t^i \leq 0$, weakly increase household inflation, *ceteris paribus*. Thus, cross-sectional variation in household inflation may, at least partly, reflect changes in consumer behaviour. Most, if not all, macroeconomic policy models currently in use by leading central banks specify models in which utility functions satisfy WARP. The empirical analysis in Section 3 shows that focusing on inflation rates for households with $\mathsf{M}_t^i \leq 0$ produces a different picture of the path of household inflation than for households with $\mathsf{M}_t^i > 0$ or for the fixed-basket measurement of inflation (CPI). The analysis in Section 4 further complicates normative (welfare) implications by showing that monetary policy directly impacts household preferences.

[5] One can express the geometric average as

$$\sqrt{\frac{\mathsf{M}_t^i(\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\mathbf{x}_t^i - \mathbf{p}_{t-1}^i(\mathbf{x}_{t-1}^i - (\mathbf{x}_t^i)(\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\mathbf{x}_t^i}{\mathbf{p}_{t-1}^i)\mathbf{x}_{t-1}^i\mathbf{p}_{t-1}^i\mathbf{x}_t^i}},$$

which demonstrates the dependence.

2.1 Comparison to standard measures

One can compare household-specific inflation to a fixed-basket measure of inflation, such as a the CPI, which holds the bundle fixed, say $\mathbf{x}_t^i = \bar{\mathbf{x}}$, $\forall t$. Cursory inspection of Equation (1) shows that $\mathsf{M}_t^i = 0$ for any fixed-basket measure. Thus, the distance between household-specific inflation, π_t^i, and a fixed-basket CPI measure, Π_t, is:

$$\pi_t^i - \Pi_t = \frac{(\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\mathbf{x}_t^i}{\mathbf{p}_{t-1}^i \mathbf{x}_{t-1}^i} - \frac{\mathsf{M}_t^i}{\mathbf{p}_{t-1}^i \mathbf{x}_{t-1}^i} - \frac{(\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\bar{\mathbf{x}}}{\mathbf{p}_{t-1}^i \bar{\mathbf{x}}}. \tag{5}$$

It follows that the distance between a household-specific inflation rate and a CPI basket depends on M_t^i in addition to bundle composition differences. An interesting question is whether the average distance in a large sample is small: does the population average of $(\pi_t^i - \Pi_t) \to 0$ as the population gets large if $\bar{\mathbf{x}}$ is the population average bundle? One can show that there is no assurance that this distance converges to zero. If N is the population size, then:

$$\frac{1}{N}\sum_{i=1}^{N}(\pi_t^i - \Pi_t) = \frac{1}{N}\sum_{i=1}^{N}\left[\frac{(\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\mathbf{x}_t^i}{\mathbf{p}_{t-1}^i \mathbf{x}_{t-1}^i} - \frac{\mathsf{M}_t^i}{\mathbf{p}_{t-1}^i \mathbf{x}_{t-1}^i}\right] - \frac{(\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\bar{\mathbf{x}}}{\mathbf{p}_{t-1}^i \bar{\mathbf{x}}}. \tag{6}$$

Whether the right-hand side of this equation is zero is an empirical question, but there is no guarantee that the population average of the first term converges to the second. Indeed, there is no reason to believe that $\sum_{i=1}^{N} \mathsf{M}_t^i/N \to 0$. Additionally, Jensen's inequality implies that the average of the first term in the summation will not in general equal the last term in the equation even if the CPI basket is the average consumption basket for the sample. It follows that there is no reason for CPI inflation to track average household-level inflation, although one cannot rule out that it may. However, such a convergence, if it did occur, would appear to be a result of random chance.

2.2 Measurement error and wealth

Echenique, Lee, and Shum (2011) propose a test to determine whether deviations from rationality are statistically significant. In their study, they find the answer is no, which suggests that perhaps the concerns over WARP raised in the previous sections are not important. One issue with the test that they propose is that it is specific to a particular form of mismeasurement that is additive in prices. The assumption of additive measurement error is potentially problematic as it implies that "true" prices could be negative.

I consider the effect of measurement error for the counterfactual prices that a household would have faced for bundles which it didn't consume. For the data used in this paper, most prices are collected at the point-of-sale by the retailer using a personal identifier card for participating retailers. For items purchased at non-participating retailers and whose price is recorded using a key by households, I assume that the price is reported accurately. While this latter assumption may appear implausible for the household-reported prices (see, e.g., Einav, Leibtag, and Nevo, 2008), in the data I am able to identify card and wand households and condition the results accordingly. However, the counterfactual data, $\mathbf{p}_t^i \mathbf{x}_{t-1}^i$ and $\mathbf{p}_{t-1}^i \mathbf{x}_t^i$, which the households experienced are not directly observed in most empirical data, including the data used in this paper. I assume that prices must be weakly positive and denote λ_1^i and λ_2^i as non-negative random vectors, with an expected value of 1, of multiplicative measurement errors for the prices. A money pump adjusted for mismeasurement can be written as:

$$\tilde{\mathsf{M}}_t^i = \mathbf{p}_{t-1}^i \mathbf{x}_{t-1}^i - (\lambda_t^i \cdot \mathbf{p}_t^i)\mathbf{x}_{t-1}^i + \mathbf{p}_t^i \mathbf{x}_t^i - (\lambda_1^i \cdot \mathbf{p}_{t-1}^i)\mathbf{x}_t^i = \mathsf{M}_t^i - (\lambda_t^i - 1)\cdot \mathbf{p}_t^i \mathbf{x}_{t-1}^i - (\lambda_{t-1}^i - 1)\cdot \mathbf{p}_{t-1}^i \mathbf{x}_t^i, \tag{7}$$

where $\lambda_t^i \cdot \mathbf{p}_t^i$ and $\lambda_{t-1}^i \cdot \mathbf{p}_{t-1}^i$ imply element-by-element multiplication.[6] For a household, $\tilde{\mathsf{M}}_t^i$ is an unbiased

[6] If this shorthand is problematic, simply consider post-multiplying prices by a matrix with λ_k^i, $k = t-1, t$, on the main diagonal and zeros elsewhere.

estimate of M_t^i if the prices and quantities are uncorrelated with the measurement errors. I discuss in the empirical section below the plausibility of this assumption for the sample I consider and its implications.

However, it is also clear from the definition of $\tilde{\mathsf{M}}_t^i$ that any measurement error that lowers counterfactual prices relative to their "true" values weakly increases the measure of money pump costs. Thus, faced with more than one plausible counter-factual price, a conservative approach to estimating money pump costs is to choose higher observed prices. It also suggests that sales or coupons for the counterfactual goods, which may also be unobserved, are likely to increase money pump costs and thus simultaneously lower household-specific inflation. Such an outcome may be difficult for policymakers to evaluate normatively, because while lower household-specific inflation may be seen to be desirable, assuming stable preferences, a simultaneous value of $\mathsf{M}_t^i > 0$ would imply that welfare interpretations may be misleading because of changing preferences.

The discussion thus far has also implicitly assumed that the bundles $\mathbf{x}_t^i$ and $\mathbf{x}_{t-1}^i$ are affordable in both periods. One might be concerned that, for example, the bundle $\mathbf{x}_{t-1}^i$ is not affordable under the price vector $\mathbf{p}_t^i$ because of a negative wealth shock to the household. Or perhaps that household expenditure is "hand-to-mouth" in the sense that the household spends its entire income in each period and the counterfactual bundle would not be affordable. To examine the impact of a binding expenditure constraint, suppose that $\mathbf{p}_t^i\mathbf{x}_{t-1}^i > w_t^i$, where w_t^i is the household's budget in period t. One can rewrite the money pump cost as:

$$\mathsf{M}_t^i = \mathbf{p}_t^i\mathbf{x}_t^i + \mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i - \mathbf{p}_t^i\mathbf{x}_{t-1}^i - \mathbf{p}_{t-1}^i\mathbf{x}_t^i.$$

The first two terms on the right-hand side are the observed expenditures and are therefore affordable to the household. Without loss of generality, assume that the household spends its available budget each period. If either (or both) counterfactual expenditure(s), $\mathbf{p}_t^i\mathbf{x}_{t-1}^i$ and $\mathbf{p}_{t-1}^i\mathbf{x}_t^i$, cost more than the observed purchases, then they are not affordable and lower M_t^i ($\mathsf{M}_t^i < 0$ if both are unaffordable). Thus, binding counterfactual affordability constraints weakly decrease money pump costs and, by implication, weakly increase household inflation given the observed bundle purchases. I note that for fixed-consumption basket measures such as the CPI, there is no role for affordability for the construction of household inflation measures whereas the money pump cost includes them.

3 Data

The data used in this paper is from the IRI Marketing Behaviorscan dataset, which contains scanner data for 31 product categories for roughly 11,000 households over a period of 11 years, 2001–2011, although not all households are present in all years.[7] The scanner data includes weekly data on the quantity of items purchased, the prices paid, and the UPC codes that identify individual products. The product categories include goods normally purchased at grocery and convenience stores (e.g., beer, razors, soft drinks). The product descriptions include information on the brand, vendor, product (e.g., Budweiser lager 355 mL), and product qualities (e.g., non-fat, 355 mL). The dataset also includes information on the UPC generation and type of system that they use to scan products. These data can be matched to the scanner data to identify items across stores and UPC generation. There are two metropolitan regions in the dataset: Eau Claire, Wisconsin, and Pittsfield, Massachusetts. In any given month, there is an average of roughly 6,000 households in the sample representing around 15,000 people.

There are two methods for BehaviourScan panelists to record consumption data. The first is to use a card, similar to a loyalty card, at participating retailers, which then records the quantities and prices paid at the point of sale. Because some stores do not participate in the survey, only a subset of consumption is observed for these households. The second method is to use a key that records their purchases from non-participating retailers (typically this is done at home) and then to enter the prices paid. Panelists can use cards or keys but all key panelists are also card panelists. Some large and significant retailers do not

[7]We thank IRI for making the data available. All estimates and analysis in the paper, based on data provided by IRI, are by the authors.

participate in the card sample, which implies that there may be a selection effect in terms of which goods are purchased by card panelists in comparison to key panelists. However, using only key panelists would reduce the sample size by roughly 90 percent and would likely increase the incidence of measurement error for prices. Thus, I include both key and card panelists in the sample. Bundle selection is problematic to the extent that it is correlated with items relatively more likely to lead to money pump costs. Thus, where feasible, I report results for key and card panelists separately. Households are in the sample for an average of roughly 9 years. In terms of physical stores, there are 798 in the sample for 2001 and 560 in 2011. Store identities are masked by IRI.

The 31 categories in the IRI data are: beer, carbonated beverages, cigarettes, coffee, cold cereals, frozen dinners, frozen pizza, hot dogs, margarine or butter, mayonnaise, milk, mustard and ketchup, peanut butter, salty snacks, soup, spaghetti sauce, sugar, yogurt, toothpaste, shampoo, and photography. I omit photography from our sample because of the significant technological change in this product category over the period of our sample (i.e., the transition to digital photography). Thus, I focus on 30 categories relating to food consumption, cleaning supplies, and regular personal care. While these items are only a subset of products that are likely purchased by households in any month, they are also unlikely to be subject to unobserved quality concerns as, for example, meats, fruits, and vegetables would be. The final sample include 4,197 different products by the L5 product classification and 25,104 products by UPC code. The model UPC item purchased is Campbell's condensed tomato soup in a 10.75 oz can. The average number of unique UPC products purchased by households in a month is 15.5, and 1.4 of these products were purchased in the previous month in the same quantity (3.1 of these products were purchased in different quantities). This simple calculation shows that, at least for the sample of households I consider, there is significant month-to-month variation in their consumption choices. The average monthly expenditure by households on products in these categories in 2001 was \$56 per month and \$68 per month in 2011. I observe approximately 11 million monthly purchases in total. The heterogeneity in households' consumption bundles and the within-household variability in purchases over time highlights how unrepresentative a fixed consumption basket may be at the UPC level for households.

Figure 1 shows the extensive margin of participating households. While there appears to be a yearly change in the number of households, 2003 was a period associated with a large number of households exiting the sample. There is also a general downward trend in the number of participating households in both Eau Claire and Pittfield although, in general, the number of participating households is almost identical across both regions in each month. In 2011, there are roughly half of the number of households in the sample as in 2001. Not all instances of having no valid observation for a household in each month is due to the household permanently leaving the sample. If a household does not meet the reporting standards set by IRI in a given year then they are dropped for that year. However, if they improve their reporting they may re-enter the sample at a later date. There are 132 months in the full sample, and 452 households in Eau Claire and 378 households in Pittfield participated in all of them.

To provide some sense of how representative the final sample is of its metropolitan area, I compare some socio-demographic characteristics between members in our sample for 2001 and the 2000 US Census. Table 1 presents the summary statistics. For the most part, the IRI sample appears broadly representative. The most notable difference is that the IRI sample over-samples home owners and also individuals lower in the income distribution. These differences appear to be explained by the IRI sample having an over representation of retired individuals as roughly one-third of the sample in 2001 has either a retired household head or a retired spouse (or both). I also note that households in our sample trend older over time because individuals age.

Although the IRI data are reported weekly, I aggregate household expenditures to the monthly frequency for two reasons. First, not all households shop on a weekly frequency so I aggregate to the monthly frequency to mitigate biases caused by infrequent purchases. Second, the focus of this paper is to examine the effect of household preferences for measures of inflation, which are typically reported at the monthly frequency.

Calculating money pump costs requires a measure of the prices for the current and past consumption bundles, which are observed, and the counterfactual prices that a household would have paid for consuming those bundles in different periods. Counterfactual prices are, as discussed previously, challenging to construct

Figure 1: Households in sample by month

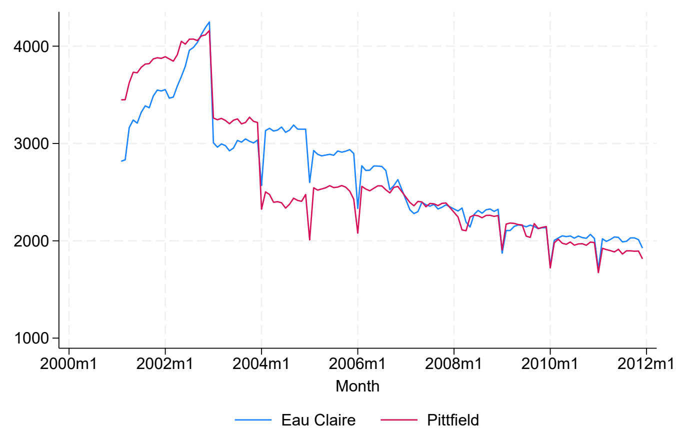

Table 1: Comparison between 2001 IRI sample and the 2000 US Census, 2001[a]

	Eau Clair, WI		Pittsfield, MA	
	Census	IRI	Census	IRI
Households	24,016	4,348	19,704	4,121
Family Income				
under $25,000	16%	30%	22%	31%
over $75,000	21%	14%	22%	14%
Households with child(ren) under 18	23%	31%	27%	25%
Average household size	2.4	2.5	2.3	2.5
Home Owners	57%	81%	61%	76%

[a] Source: U.S. Census Bureau, Census 2000 Summary File 1.

if consumption bundles change because the IRI data does not include prices for goods not purchased in a given period (i.e. the IRI data does not include shelf prices). Thus, counterfactual prices are unobserved and the prices that could have been paid by consumers for a particular good may depend on what stores they visit in a month, when they visit, the current stock of the store (e.g., if a store runs out of 2 L bottles of Coca-Cola so only 591 mL bottles are available).[8]

To construct the counterfactual prices, I proceed in two steps. First, I use the full BehaviourScan sample by region to calculate the median price paid for the UPC item by all households in the sample. If at least one household has purchased the UPC item, then this step provides the counterfactual value.[9] If, however, no household in the sample has purchased the UPC item, I proceed to step 2 and use IRI's store-level data for Eau Claire and Pittfield to calculate the average price paid for that item according to its description.[10] To calculate the average price paid, for each store I divide the total dollar sales of the item by the quantity of units purchased. Using steps 1 and 2 generates counterfactual prices for roughly 97% of the roughly 11 million goods in the sample. I note that there are two periods when counterfactual prices cannot be constructed because of differences in how the data are collected. UPC codes change in January, 2007, and January, 2008, which implies that items cannot be linked at the UPC level for the periods December 2006 to January 2007 and December 2007 to January 2008. I drop these periods from the data whenever item-level price comparisons are required.

Although counterfactual prices for goods not consumed by the household are unobserved, it is possible to compare the constructed counterfactual prices with the prices actually paid by the household for goods purchased in both periods (for which the prices are observed). For these repeated purchases, I calculate the ratio of the constructed counterfactual prices to the price actually paid. Because there is typically more than one possible counterfactual price (e.g., there can be more than one other household or store that purchased the item in the counterfactual period), I consider four different definitions of the counterfactual price: the median, the mean, the minimum price, and the maximum price. Table 2 summarizes the ratio of each of these four prices to the price actually paid. Both the median and the mean counterfactual prices appear to be broadly representative of the price actually paid, as the median and mean value of these ratios are essentially 1 (the medians are both 1 and the means are 1.01). The median counterfactual price is with 3 percent of the price actually paid for over half of the roughly 2 million repeat purchases observed. And it is within 15 percent for almost 80 percent of the transactions.[11] The dispersion of the mean counterfactual price ratio appears slightly larger. For both the median and mean, the difference between the actual price paid and the counterfactual predicted price appears symmetrically distributed, which is suggestive evidence that the measurement error is uncorrelated with the price. In comparison, both the minimum and maximum counterfactual prices appear to skew the distributions and have mean and median ratios different from 1 by at least 10 percentage points. For the analysis that follows, I use the median counterfactual price to calculate money pump costs and related statistics.

4 Consumption and expenditure patterns

4.1 Key vs. card panelist behaviour

I calculate monthly average expenditure separately for key and card panelists and graph the resulting time series (see Figure 2a). Because card expenditures are, at least weakly, a subset of key expenditures (which

[8] An additional issue is that the prices we observe may not be the actual prices paid if consumers have a discount coupon issued by the vendor. While IRI records retailer discounts, such as price reductions by a grocery store, IRI does not observe coupons by, for example, Kellogg's for a cereal brand.

[9] For the data analysis in this paper, I use Stata. To construct the median value of an even number of observations, Stata calculates the average value of the central two order statistics.

[10] In the IRI data, the product descriptions are at the L9 level, which is the same level as the UPC codes. I discuss next cases when the mapping from L9 description to UPC classification is problematic.

[11] The distribution of the median counterfactual prices relative to the actual price paid for a repeated purchase is essentially identical for both key and card panelists, respectively. This decomposition is available upon request.

Table 2: Ratio of counterfactual prices and actual prices paid for repeat purchases

Counterfactual price used	Percentiles					
	10th	25th	median	75th	90th	mean
Median	0.84	0.97	1	1.01	1.15	1.01
Mean	0.85	0.94	1	1.05	1.19	1.01
Minimum	0.60	0.74	0.90	1	1	0.85
Maximum	1	1	1.10	1.29	1.60	1.21

Median refers to the median price observed for that item in the counterfactual data; Mean refers to the mean price observed; Minimum refers to the minimum observed price; and Maximum refers to the maximum observed price. There are 2,103,705 items that are classified as repeat purchases at the UPC level out of the 11,218,520 purchases observed in the data.

also include card purchases), one should expect that average key panelist expenditure is higher than card panelist expenditure. This appears to be generally true in the figure. This relation need not hold strictly, however, because key and card panelists may have different consumption levels for a variety of reasons, such as wealth, household composition, or the quality of items purchased. Nevertheless, broadly speaking, key and card panelists appear to have nearly identical consumption trends over the sample period.

Even if key and card panelists are similar in the level of expenditure, they may differ in terms of the number of items in their consumption basket. Figure 2b plots the time series of the average number of unique UPC items purchased in each month by key and card panelists. There appears to very little difference in terms of the consumption bundle size between key and card panelists, although key panelists do typically purchase slightly more goods than card panelists. Over the sample, the average number of unique items purchased is roughly stable, though there is some evidence that fewer goods are purchased after around mid-2006.

Figure 2: Key and card panelist average expenditure and average unique items purchased, by month

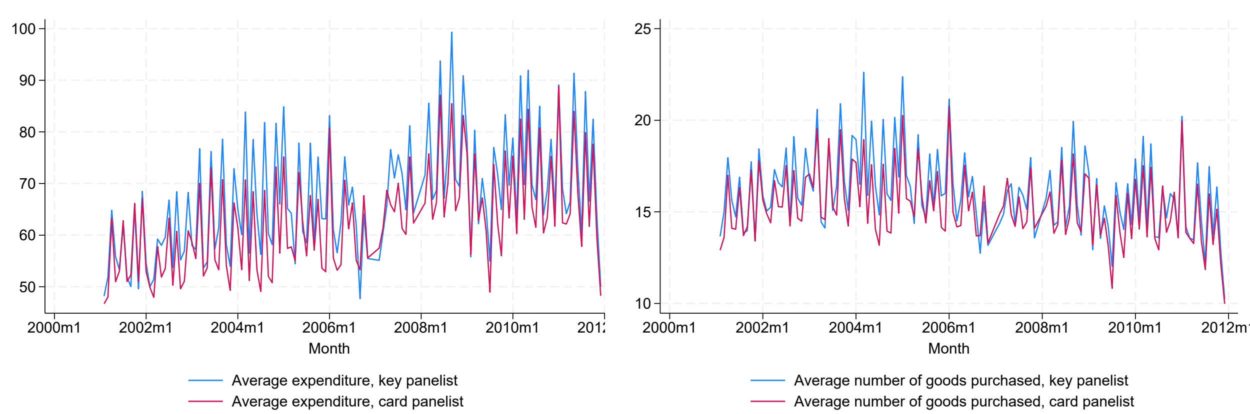

(a) Average expenditure

(b) Average unique items purchased

Although there appears to be little difference in the number of unique items purchased between key and card panelists, it remains possible that there are differences in the number of repeat purchases. Figure 3 plots the average number of repeat purchases, defined as purchases of the same units of the same unique item in adjacent months, for key and card panelists. The time-series behaviour of both key and card panelists is similar, however there is a noticeable decline in the number of repeat purchases after 2007. This suggests that a fixed-basket measure of prices for these households is relatively less representative of the inflation pressure on households after 2007.

Figure 3: Key and card panelist average repeat unique items purchased, by month

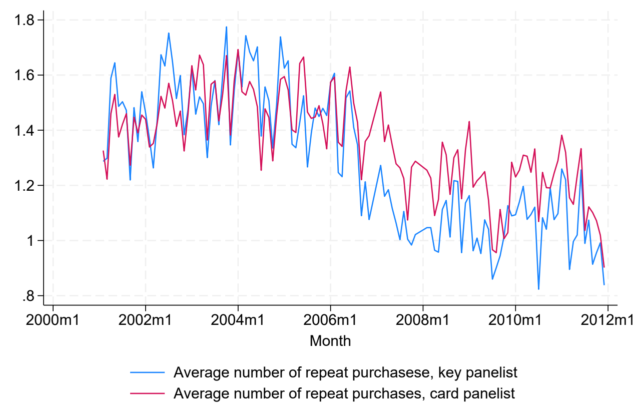

Figures 4a and 4b plot the relative inflationary pressures faced by key and card households for repeat purchases versus all purchases. For both key and card panelists, the average inflation for repeated purchases is lower than for all items. The decline in repeat purchases will mechanically increase household inflation for all goods, because average inflation for repeated purchases is relatively lower than for other goods. This effect appears to be present in the figures, as there is an increase in average inflation around 2007 for key, and possibly card, panelists, which corresponds with the decline in the average number of repeat purchases around 2007 observed in Figure 3.

Figure 4: Key and card panelist average month-to-month inflation for repeat purchases and all items

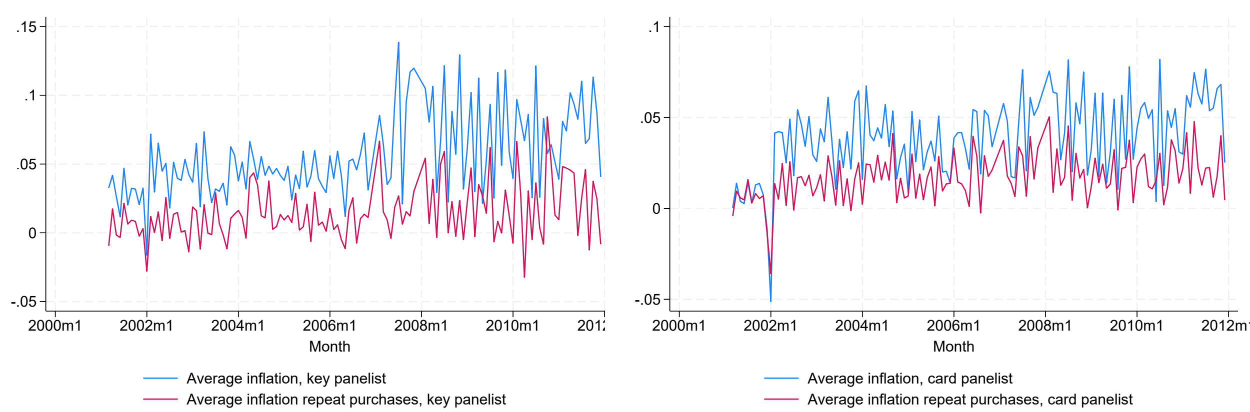

(a) Key panelists

(b) Card panelists

4.2 Money pump costs and household inflation

While the dynamics of household inflation are potentially interesting of their own right, one central point of this paper is that household preferences affect inflation, as shown in Equation 2. Figures 5a and 5b plot the decomposition of average household inflation given by Equation 2 for the key and card panelists. The red line in both figures is the difference between the nominal and real expenditure growth terms in Equation

2: $\frac{\mathbf{p}_t^i\mathbf{x}_t^i}{\mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i} - \frac{\mathbf{p}_{t-1}^i\mathbf{x}_t^i}{\mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i}$. The blue lines are the average household month-to-month inflation. Finally, the green line in both figures is the contribution of the average money pump cost: $-\frac{\mathsf{M}_t^i}{\mathbf{p}_{t-1}^i\mathbf{x}_{t-1}^i}$. For both key and card panelists, money pump costs raise household inflation. Indeed, average household inflation would be negative almost entirely throughout the sample period without the contribution of money pump costs.[12]

Figure 5: Key and card panelist average month-to-month inflation decomposition

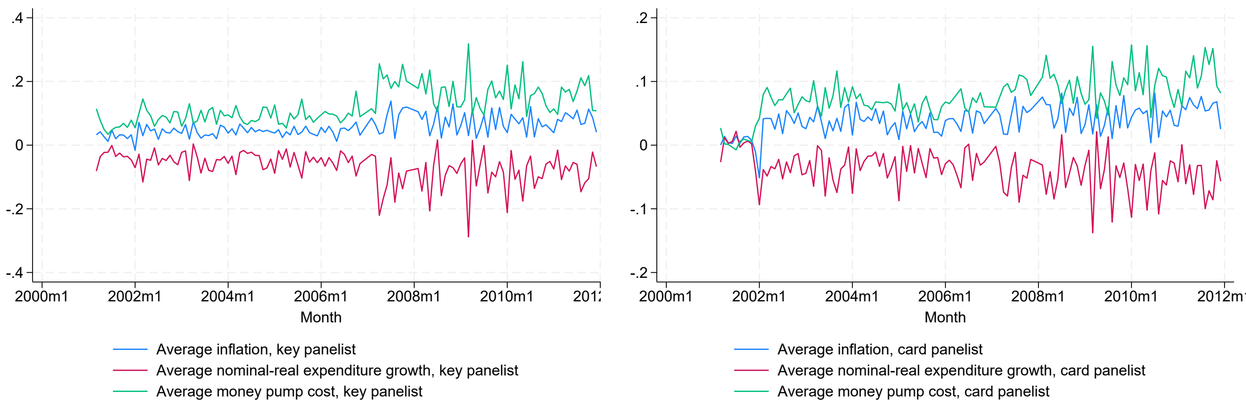

(a) Key panelists

(b) Card panelists

The average money pump cost is not, however, a revealing measure of household heterogeneity with respect to money pump costs nor how dispersion in inflation rates may reflect these costs. The decomposition of inflation implies that positive (negative) money pump costs decrease (increase) inflation. To illustrate the empirical relevance of money pump costs at the household level, Figure 6a plots the average fraction of key and card households in each month who have $\mathsf{M}_t^i > 0$. While $\mathsf{M}_t^i > 0$ appears relatively common for both key and card panelists, the fraction of card panelists with positive money pump costs is higher than that of key panelists by roughly 5 to 10 percentage points.[13] Figure 6b plots average household inflation for panelists with $\mathsf{M}_t^i \leq 0$ and $\mathsf{M}_t^i > 0$. Over the sample period, households with money pump costs $\mathsf{M}_t^i \leq 0$ experienced relatively high month-to-month inflation rates, averaging slightly under 10 percentage points for the 30 product categories examined in this paper. Conversely, households with money pump costs $\mathsf{M}_t^i > 0$ experienced deflation of roughly 4 percentage points on average for the same product categories over the same sample period. This evidence suggests that household preferences, as reflected by money pump costs, are closely tied to the inflation experienced by the household.

[12] One may be concerned that the choice of using the median counterfactual price is driving the qualitative results. Figure 11a plots the average fraction of (both key and card) households with $\mathsf{M}_t^i > 0$ by three different measures of the counterfactual prices: the median, minimum, and maximum. Although there are obvious differences in the proportion of panelists with positive money pump costs according to which the counterfactual price is used, all indicate that at least some households do not satisfy WARP. It is perhaps instructive to recall that values of $\mathsf{M} < 0$, increase household inflation *ceteris paribus*, so opting for a counterfactual price that implies fewer WARP violations will necessarily increase average household inflation. Figure 11b illustrates this point.

[13] This is perhaps not entirely unexpected since the consumption by card panelists is from fewer stores, which may reflect less search effort on the part of these households. As suggested by Echenique, Lee, and Shum (2011), one interpretation of M_t^i is as the amount of money that could be extracted by a devious arbitrageur, in this case the store. Thus, if stores are aware of less search effort by households, they may exploit this to earn additional profits.

Figure 6: WARP violations

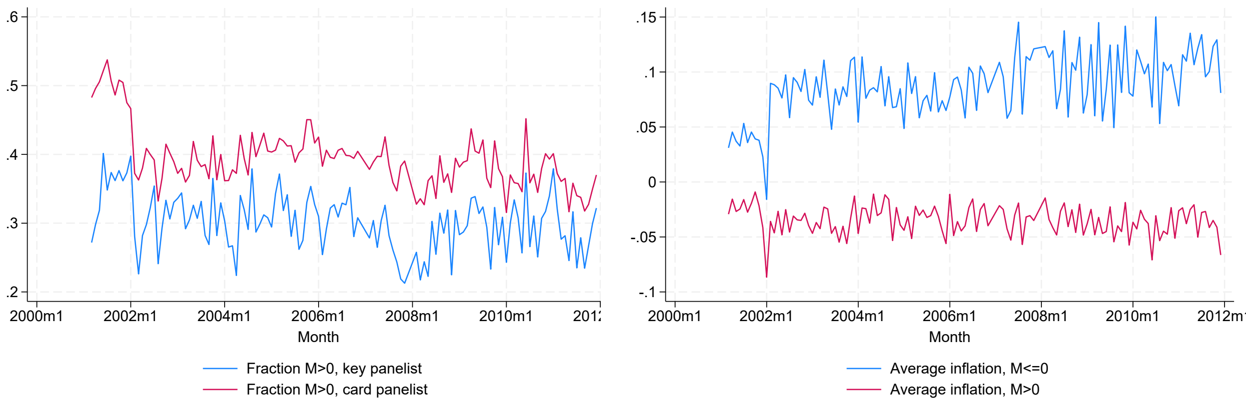

(a) Key and card panelists

(b) Average household inflation, month-to-month

4.3 Individual heterogeneity

I next turn to individual time-series heterogeneity: How often do households have positive money pump costs? For each household in the sample, I calculate the proportion of months in the sample that the household has $\mathsf{M}_t^i > 0$. Table 3 presents summary statistics by card and key panelists. On average, card and key panelists are qualitatively similar and tend to have positive money pump costs 30 percent of the time they are in the sample. For both key and card panelists, there are some households that never have positive money pump costs (roughly 5 percent of the card and 1 percent of the key panelists). There are also no card or key households that always have positive money pump costs. Although the distributions are qualitatively similar, it is generally the case that key panelists have positive money pump costs less frequently than card panelists.

Table 3: Proportion of months in sample with $M_t^i > 0$, by household

		Percentiles						
	min	10th	25th	median	75th	90th	max	mean
Card panelist	0	0.085	0.188	0.314	0.449	0.579	0.953	0.325
Key panelist	0	0.068	0.142	0.241	0.364	0.477	0.792	0.263

There are 10,757 card panelist households and 923 key panelist households in the sample. "min" and "max" refer to the minimum and maximum proportions across households.

To illustrate how much money pump costs could matter for household inflation, I run a simple regression to determine the correlation between money pump costs and household inflation:

$$\pi_t^i = \alpha_i + \delta_t + \beta \mathsf{I}(\mathsf{M}_t^i > 0) + \gamma KS_t^i + \mu KS_t^i \times \mathsf{I}(M_t^i > 0) + e_t^i \tag{8}$$

where α_i is an individual household fixed effect, δ_t is a time dummy variable, β is the effect on household inflation if $\mathsf{M}_t^i > 0$, KS_t^i is an indicator variable for a key panelist (some households switch types, so this is not subsumed by the fixed effect), γ measures the effect of panelist type, and μ captures the interaction between panelist type and M_t^i. If this latter coefficient is statistically significantly different from zero, then this suggests that positive money pump costs have different effects on household inflation by panelist type. β is the coefficient of interest, as this measures the effect of the extensive margin of money pump costs for household inflation. Table 4 presents the results. As the estimates in the table indicate, household inflation is roughly 10 percentage points lower when households have positive money pump costs. Given that households tend to have $\mathsf{M}_t^i > 0$ roughly 30 percent of the time, then the estimate $\hat{\beta}$ suggests that it is not uncommon for households to face 10 percentage point swings in their household inflation rates because of changes in their preferences. This regression evidence is, however, not casual, as there may be an underlying factor linking money pump costs and household inflation.

Table 4: WARP and household inflation

	Estimate	Std err	t statistic
$\hat{\beta}$	-0.102	0.0003	-316
$\hat{\gamma}$	-0.008	0.0016	-4.85
$\hat{\mu}$	0.0003	0.0010	0.30
Observations		637,942	
R^2		0.26	

There are 11,217 households in the full sample and 126 months. α_i and δ_t are included in the regression specification.

5 Monetary policy, preferences, and household inflation

As noted in the introduction, a central question in macroeconomics is to understand how monetary policy affects inflation. Many central banks have an explicit inflation target, and canonical New Keynesian theories and models embed mechanisms through which changes in monetary policy lead to dynamic adjustments in inflation. The channels in these models are generally consistent with the decomposition of household inflation in this paper, given by Equation (2). The first term in Equation (2), nominal expenditure growth, captures pricing decisions by firms for consumption goods; the second term, the Laspeyres quantity index, reflects household real spending that is typically related to real incomes (employment); and the third term, money pump cost, reflects changes in relative demand for particular goods. This latter channel is emphasized in recent models that analyze the macroeconomic effects of the pandemic (e.g., Baqaee and Farhi (2022)), although these authors do not consider a model with monetary policy. To understand how monetary policy affects household inflation, I next analyze its effect on each component in Equation (2) separately.

I focus on the average household inflation rate as the monetary policy target variable and denote $\bar{a}_t$ as the period t average of a given variable a (i.e., over the $i \in N_t$ households in the sample in that period). The decomposition of household inflation implies that:

$$
\begin{aligned}
\bar{\pi}_t &= \bar{E}_t - \bar{R}_t - \bar{M}_t \qquad \text{where} \\
\bar{E}_t &= \frac{1}{N_t}\sum_{i \in N_t} E_t^i = \frac{1}{N_t}\sum_{i \in N_t} \frac{\mathbf{p}_t^i \mathbf{x}_t^i}{\mathbf{p}_{t-1}^i \mathbf{x}_{t-1}^i} \\
\bar{R}_t &= \frac{1}{N_t}\sum_{i \in N_t} R_t^i = \frac{1}{N_t}\sum_{i \in N_t} \frac{\mathbf{p}_{t-1}^i \mathbf{x}_t^i}{\mathbf{p}_{t-1}^i \mathbf{x}_{t-1}^i}, \quad \text{and} \\
\bar{M}_t &= \frac{1}{N_t}\sum_{i \in N_t} M_t^i = \frac{1}{N_t}\sum_{i \in N_t} \frac{\mathsf{M}_t^i}{\mathbf{p}_{t-1}^i \mathbf{x}_{t-1}^i}.
\end{aligned} \tag{9}
$$

One well-known challenge with estimating the effect of monetary policy on macroeconomic aggregates is that it is difficult to isolate the counterfactual outcome that would have occurred absent a policy change. This challenge is often resolved by an assumption that the macroeconomic counterfactual can be estimated by an autoregressive process, such as a vector autoregression (VAR). In addition, because the time dimension of many macroeconomic data series is relatively short, the macroeconometric forecasting literature has also examined dimension reduction techniques; see Bai and Ng (2002), Stock and Watson (2002), and Stock and Watson (2012).

In this section, I borrow intuition from this literature, in particular, the three pass regression filter proposed by Kelly and Pruitt (2015), to estimate the impact of monetary policy for average household inflation without relying on VAR dynamics to estimate the counterfactual.[14] For each $\Theta = \{E, R, M\}$, the first stage of the three pass regression filter is a time-series regression for each i:

$$\Theta_t^i = \alpha_{i,0}^{\Theta} + S_t' \beta_i^{\Theta} + e_{i,t}^{\Theta} \tag{10}$$

where S_t is the unobserved monetary policy factors (here the monetary policy changes). (As discussed next, I choose the federal funds rate and policy news shocks of Nakamura and Steinsson (2018) as proxies.) The second stage of the three pass regression estimates the unobserved factors using a cross-section regression for each period t:

$$\Theta_t^i = \alpha_{0,t}^{\Theta} + \hat{\beta}_i^{\Theta'} F_t^{\Theta} + \ddot{e}_{i,t}^{\Theta} \tag{11}$$

where $\hat{\beta_i^{\Theta}}$ is the coefficient estimates from the first stage regression and F_t is the factors (note that with two shocks, F_t is a 2×1 vector). The second stage regressions also include a constant for each time period, which capture the effects of other variables that may be coincident to the monetary policy shocks. Kelly and Pruitt (2015) show that these factors converge to the infeasible best linear forecast in a predictive regression, such as:

$$\bar{\pi}_{t+h} = \delta_0 + \hat{F}_t^{\Theta'} \gamma_h^{\Theta} + u_{i,t+h}. \tag{12}$$

This predictive regression is similar to a local projection specification (Jordà (2005)). Before returning to the issue of the dynamic effect of monetary policy on inflation, consider the second stage regression. It implies that:

$$\bar{\Theta}_t = \alpha_{0,t}^{\Theta} + \bar{\hat{\beta}}_i^{\Theta'} \hat{F}_t^{\Theta}. \tag{13}$$

Because the inflation decomposition is an identity, this implies:

$$\bar{\pi}_t = \alpha_{0,t}^{E} + \bar{\hat{\beta}}_i^{E'} \hat{F}_t^{E} + \alpha_{0,t}^{R} + \bar{\hat{\beta}}_i^{R'} \hat{F}_t^{R} + \alpha_{0,t}^{M} + \bar{\hat{\beta}}_i^{M'} \hat{F}_t^{M} + o_t, \tag{14}$$

which decomposes the effect of monetary policy on inflation and provides an estimated counterfactual for $\bar{\pi}_t$ in the absence of monetary policy shocks by setting the relevant factors to zero. The term o_t captures all non-monetary factors that feed into inflation that are not captured by the fixed effects from the individual regressions, $\alpha_{0,t}^{E}, \alpha_{0,t}^{R}$, and $\alpha_{0,t}^{M}$. It is worth noting that the decomposition identification in Equation (14) is based on the assumed factor structure for each component instead of a VAR and so is different than the identification strategy typically used in the macroeconometrics literature on monetary policy.

Returning to the dynamic effects of monetary policy shocks on inflation, Equation (13) also assumes a factor structure for inflation. This assumption is trivial for $h = 0$ as it is implied by the factor structure assumption for the separate components. However, for $h > 0$, interpreting the γ_h^{Θ} as the impulse responses from a monetary policy shock does require an assumption that the data-generating process underlying $\bar{\pi}$ is a VAR(p) that depends at least in part on the monetary policy factors $\hat{F}_t^{\Theta}$.

[14] One relative advantage of the three pass regression filter is that it easily handles missing observations which are a feature of the household-level consumption data.

5.1 Monetary policy shocks

Since monetary policy decisions presumably depend on the state of the economy, causal identification of monetary policy effects requires either an instrumental variable or an identified unexpected monetary policy change—a monetary policy shock. There appear to be two main policy levers available to central banks: changes to the overnight rate of interest and announcements about future policy actions (sometimes referred to as "forward guidance"). Identifying unexpected monetary policy changes is generally done using one of several methods. One method is to use a structural VAR to overcome endogeneous monetary policy changes by controlling for the factors that might lead to those changes using restrictions imposed on that VAR (e.g., Christiano, Eichenbaum, and Evans (1999) and Bernanke, Boivin, and Eliasz (2005)). It is, however, difficult to map those restrictions to the (relatively) model agnostic approach taken in this paper. Another method to identify monetary policy shocks, the narrative approach, was proposed by Romer and Romer (2004) using the Federal Reserve's internal forecasts for the federal funds rate and minutes from FOMC meetings to construct a measure of federal funds rate changes orthogonal to the policymakers' information.

A third method to identifying monetary policy shocks is to use high frequency changes in market prices around monetary policy announcements to identify the unexpected component. This approach assumes market efficiency—that interest rates just prior to a monetary policy announcement reflect the expectation of that announcement. Cook and Hahn (1989) and Kuttner (2001) are early pioneers of this approach, and Cochrane and Piazzesi (2002) demonstrate how this approach can help disentangle coincidental shocks from the monetary policy shock, which is challenging using VAR methods. A recent contribution to this second approach is Nakamura and Steinsson (2018), who focus on bond price movements in a narrow, 30-minute window around US Federal Reserve announcements. Nakamura and Steinsson (2018) also identify two monetary policy shocks: a federal funds rate shock (ffr) and a policy news shock (policy). The federal funds rate shock measures the change in the federal funds rate in a 30-minute window around the announcement of a FOMC meeting. The policy news shock is "a composite measure of changes in interest rates at different maturities spanning the first year of the term structure." Importantly, the policy news shock captures the effect of forward guidance. I use these shocks as the proxies S_t in the first stage regression for the three pass regression filter, Equation (10). For the second stage regression, I also restrict the sample to include only months during which there was a FOMC meeting.

Table 5 presents some summary statistics for the estimated factors. As noted by Nakamura and Steinsson (2018), the federal funds rate and policy news shocks are relatively small in terms of their magnitude, with standard deviations of only a few basis points. The factors estimated from these shocks are likewise relatively small in magnitude, which does raise concerns about possible power problems when used as regressors. However, it remains useful to analyze their impact on household inflation as they represent what is effectively the non-mechanical aspect of monetary policy decision making.

I begin by examining the effect of the monetary factors for contemporaneous inflation as presented in Equation (14). Traditional economic theory suggests that monetary shocks do not contribute contemporaneously to household inflation, and the evidence in Figure 7 corroborates this view. The contemporaneous impact of the monetary factors contributes very little to household inflation in level terms and is essentially uncorrelated with it (unconditional correlation coefficient of 0.016). This evidence suggests that monetary policy changes, if they have any effect at all, must operate with a lag, which is consistent with theory.

Table 5: Summary statistics for estimated factors, shocks, and inflation targets

variable	mean	std dev	min	max
$\hat{F}_t^E(ffr)$	-0.006	0.069	-0.441	0.149
$\hat{F}_t^E(policy)$	-0.002	0.048	-0.238	0.138
$\hat{F}_t^R(ffr)$	-0.005	0.068	-0.430	0.152
$\hat{F}_t^R(policy)$	-0.002	0.048	-0.233	0.140
$\hat{F}_t^{ER}(ffr)$	-0.002	0.063	-0.269	0.193
$\hat{F}_t^{ER}(policy)$	0.005	0.047	-0.153	0.152
$\hat{F}_t^M(ffr)$	0.000	0.061	-0.249	0.200
$\hat{F}_t^M(policy)$	0.010	0.047	-0.144	0.149
NS ffr	-0.007	0.056	-0.413	0.125
NS policy news	-0.002	0.039	-0.243	0.099
$\bar{\pi}_t$	0.041	0.022	-0.048	0.088
$\iota(\pi)_t$	0.127	0.029	0.048	0.186
Observations		126		

Notes: ffr refers to the federal funds rate shock; policy refers to the policy news shock; NS refers to the original Nakamura and Steinsson (2018) federal funds and policy news shocks; $\bar{\pi}_t$ is the cross-sectional average household inflation in period t; and $\iota(\pi)_t$ is the cross-sectional inter-quartile range of household inflation in period t. There are 126 observations for each variable.

Figure 7: Contribution of monetary shocks to contemporaneous average inflation, by month

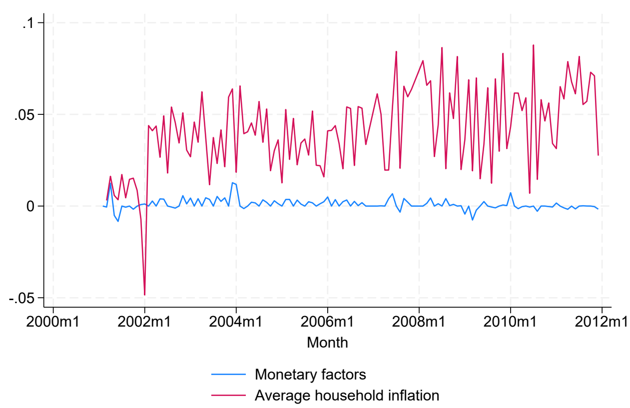

5.2 The dynamics of monetary policy

To investigate the effect of monetary policy for the dynamics of household inflation, I estimate a Jordà (2005) local projection regression using lags of household inflation to control for serial correlation as recommended by Montiel Olea and Plagborg-Møller (2021). For expositional clarity, I also combine the factors for nominal expenditure and real expenditure because, in level terms, these factors measure gross growth rates and are orders of magnitude larger than inflation or money pump changes and also have offsetting effects.[15] The local projection specification is:

[15] Combining the nominal and real factors does not qualitatively change the results. The disaggregated results are available upon request.

$$\bar{\pi}_{t+h} = \lambda_m + \hat{F}_t^{ER}\mu_h^{ER} + \hat{F}_t^{M}\mu_h^{M} + \sum_{i=1}^{13}\rho_i\bar{\pi}_{t-i} + \gamma_{t+h}FOMC_t + \epsilon_{t+h} \tag{15}$$

where λ_m is a vector of monthly dummy variables to capture seasonal patterns in inflation rates, $\hat{F}_t^{ER} = \hat{F}_t^{E} - \hat{F}_t^{R}$ is the combined nominal and real factor, $FOMC_t$ is a dummy variable for whether there was a FOMC meeting in that month, and $h = 1, ..., 24$ is the local projection horizon. I choose 12 lags of the dependent variable to control for serial correlation in inflation (this choice appears conservative, as the regression residuals do not display any evidence of serial correlation after including 8 lags). Figure 8a plots the estimated coefficients, $\hat{\mu}_h^{ER}$ for the Nakamura and Steinsson (2018) federal funds rate shock for the nominal and real factor from the local projection. The estimated impulse responses show almost no statistically significant effect from monetary policy for average household inflation for a horizon of up to 12 months. There is evidence of a statistically significant federal funds rate effect operating through the nominal and real factors in alternating months between 14 and 22 months, although this effect appears offset by a similar significance for the coefficients on the federal funds rate shock operating through the money pump costs; see Figure 8c. The evidence suggests that the effects of federal funds rate shocks on inflation are nuanced and will therefore depend on the magnitudes of the nominal and real factor and the money pump factor. Since the money pump factor is, on average, roughly twice as large (see Table 5), this suggests that the impact from a positive federal funds rate shock was, for the sample considered, to modestly lower inflation over the period from 12 to 24 months. Ignoring the effects of the money pump costs, federal funds rate shocks would modestly increase inflation over the same period, which is consistent with the well-known "price puzzle"; see, for example, Coibion (2012), Ramey (2016), and Miranda-Agrippino and Ricco (2021). Finally, I find no evidence of a statistically significantly effect from policy news shocks operating either through the nominal and real factor or the money pump factor; see Figures 8b and d.

5.3 Inflation inequality and monetary policy

Although I find only modest evidence that monetary policy affects average household inflation through its effects on expenditure or household behaviour, it remains possible that monetary policy has distributional effects for household inflation. For instance, it is plausible to think that household attentiveness to monetary policy is heterogeneous or that differences in household consumption baskets may lead to different household consumption responses (Lauper and Mangiante (2021)). To examine the effect of monetary policy on the distribution of household inflation, I calculate the inter-quartile range of household inflation in each period t, $\iota(\pi)_t$, and estimate a local projection similar to Equation (15):

$$\iota(\pi)_{t+h} = \mathsf{l}_m + \hat{F}_t^{ER}\mathsf{m}_h^{ER} + \hat{F}_t^{M}\mathsf{m}_h^{M} + \sum_{i=1}^{13}\mathsf{p}_i\iota(\pi)_{t-i} + \mathsf{g}_{t+h}FOMC_t + \varepsilon_{t+h} \tag{16}$$

where I have replaced the coefficients from Equation (15) with their alphabetic counterparts for clarity.

Figure 9 plots the impulse responses for the nominal and real and the money pump factors for the federal funds and policy news shocks. The inter-quartile range does not identify which households are affected by a factor (e.g., a negative coefficient would imply that the range is narrowing but whether this is because of an effect on households at the top or at the bottom of the distribution is uncertain). However, the impulse responses suggest that the federal funds rate and the policy news shocks operate at different horizons. Focusing on Figures 9b and d, it appears that policy news shocks affect the distribution of household inflation for the first 12 months of the horizon. However, the policy news shocks affect households through two different, and offsetting channels, similar to the response of average household inflation to federal funds rate shocks. The policy news shocks operating through the nominal-real factors increase the inter-quartile range of inflation, but the same shocks operating through the money pump costs lower the inter-quartile range. Thus, absent money pump costs, monetary policy news shocks increase the dispersion of inflation

Figure 8: Local projection estimates for average household inflation

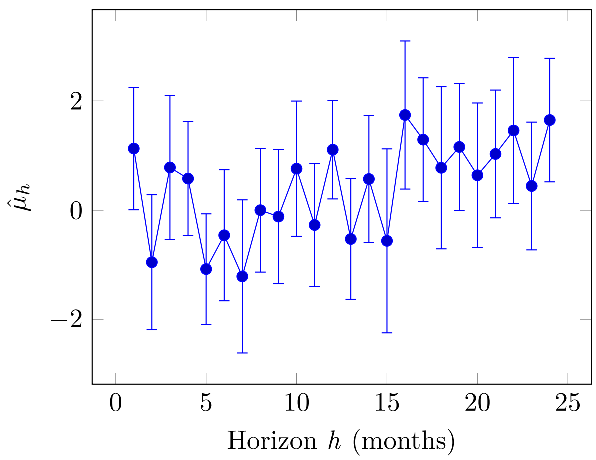

(a) Nominal - real (ER), federal funds factor

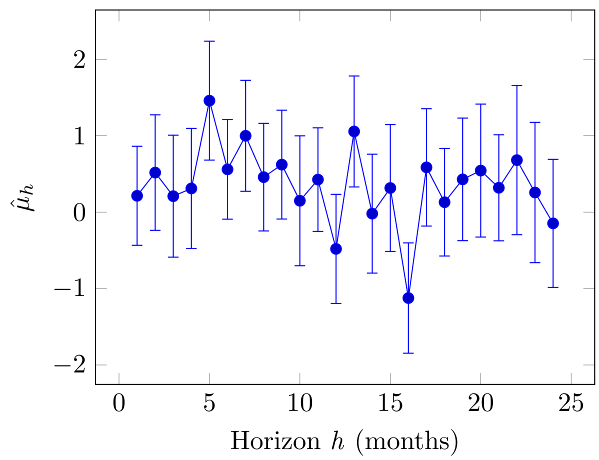

(b) Nominal - real (ER), policy news factor

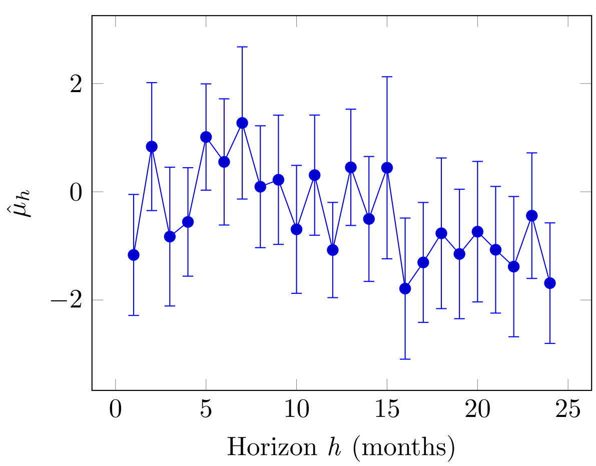

(c) Money pump (M), federal funds factor

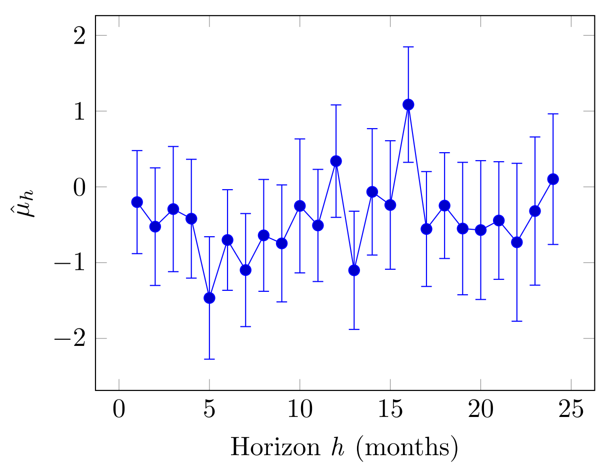

(d) Money pump (M), policy news factor

Notes: The figures show the estimated coefficients $\hat{\mu}_h$. Markers represent point estimates; whiskers indicate 95% confidence intervals using robust standard errors. The sample period underlying the estimates is monthly, 2001 to 2011. Observations for each regression range between 105 and 80 (recall 4 months are dropped from the sample because of sampling changes) for households. At the same time, monetary policy news shocks decrease the dispersion of inflation because of their influence on preferences. The overall effect of monetary policy announcements for the inter-quartile range of inflation thus depends on the relative magnitude of the factors themselves—since the money pump factor is roughly twice as large as the nominal-real factor, then the coefficient estimates suggest that the policy news shocks tend to lower inflation dispersion for the sample considered here. Because the average household inflation rate is unaffected by the policy news shocks over the same horizon, this suggests that the decline in dispersion is symmetric around average inflation. This is suggestive evidence that policy news shocks may be focal for households.

After the 12-month horizon, the estimated impulse responses from the policy news shocks are generally not statistically significantly different from zero for either the nominal-real or money pump factors. However, at this point, the estimated impulse responses for the federal funds rate shocks behave similarly to how they did for average household inflation, alternating significance in the months between 12 and 24. The estimates of the federal funds shock for the nominal-real factor is positive, which suggests that monetary tightening (easing) increases (decreases) inter-quartile inflation through this channel. However, the estimates of the federal funds shock for the money pump factor are negative in the same period, which implies that monetary tightening (easing) increases (decreases) inter-quartile inflation through this channel. Thus, similarly to the impulse responses for average household inflation, the federal funds rate shocks have offsetting effects on monetary policy through the nominal-real and money pump channels, and the importance of each for overall dispersion is determined by the size of the respective factors. Because the money pump factor is larger, this suggests that federal funds rate shocks lower dispersion, although the changes in average inflation noted above suggest that this decline is not symmetric around average inflation.

5.4 Cumulative multipliers

The evidence presented in the preceding sections suggests that policy news shocks and federal funds shocks operate at different horizons for household inflation. However, the estimated effects are, perhaps, a little difficult to visualize at the monthly frequency because of the tendency of the inflation series to revert. To better assess the effects of the shocks for inflation, I estimate cumulative multipliers as in Ramey and Zubairy (2018). For each period, I construct cumulative 1–12 month and 13–24 month changes in average household inflation and its inter-quartile range and estimate:

$$\sum_{l}^{u} \bar{\pi}_{t+h} = \ddot{\lambda}_m + \hat{F}_t^{ER}\ddot{\mu}^{ER} + \hat{F}_t^{M}\ddot{\mu}^{M} + \sum_{i=1}^{13} \ddot{\rho}_i \bar{\pi}_{t-i} + \ddot{\gamma}_{t+h} FOMC_t + \upsilon_{t+h}, \tag{17}$$

and

$$\sum_{l}^{u} \iota(\pi)_{t+h} = \ddot{\mathsf{l}}_m + \hat{F}_t^{ER}\ddot{\mathsf{m}}_h^{ER} + \hat{F}_t^{M}\ddot{\mathsf{m}}_h^{M} + \sum_{i=1}^{13} \ddot{\mathsf{p}}_i \iota(\pi)_{t-i} + \ddot{\mathsf{g}}_{t+h} FOMC_t + \vartheta_{t+h} \tag{18}$$

where $l = \{1, 13\}$ and $u = \{12, 24\}$ are the lower and upper bounds of the cumulative windows, respectively. The estimates are presented in Table 6. For average inflation, the estimated coefficients for the policy news shocks and the federal funds shocks are significant for the 1–12 and 13–24 month cumulative windows respectively and, similar to the monthly results, are of opposite sign and roughly the same magnitude. However, because the values of $\hat{F}_t^{ER}$ and $\hat{F}_t^{M}$ are also of different magnitudes, the net effect on cumulative inflation is non zero. Figure 10a plots the contribution to a one-year cumulative horizon of inflation for the policy news shocks and federal funds shocks in the reported month for the 1–12 and 13–24 cumulative windows, respectively.[16] The net effect on cumulative inflation for both shocks is generally negative over the

[16] How to report the time axis of this chart is somewhat arbitrary, so I have linked the month to the timing of the monetary shock rather than its impact. Thus the policy news effect reported for, say, January 2008 is the one-year cumulative effect in January 2009. Similarly, the federal funds effect reported for January 2008 would be the one-year cumulative effect in January

Figure 9: Local projection estimates for inter-quartile household inflation

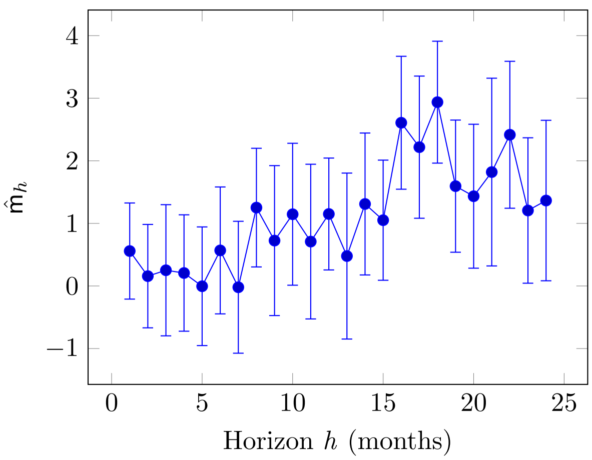

(a) Nominal - real (ER), federal funds factor

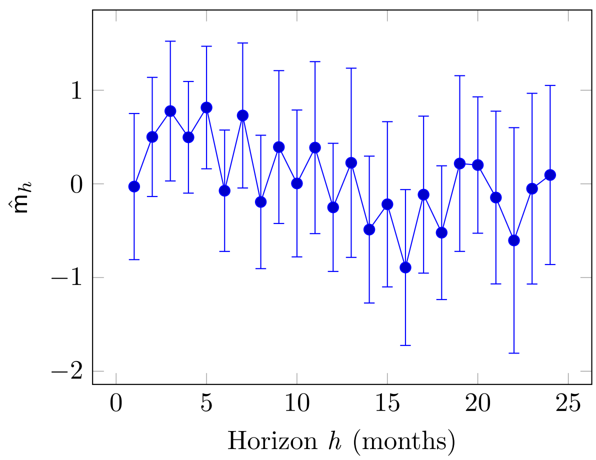

(b) Nominal - real (ER), policy news factor

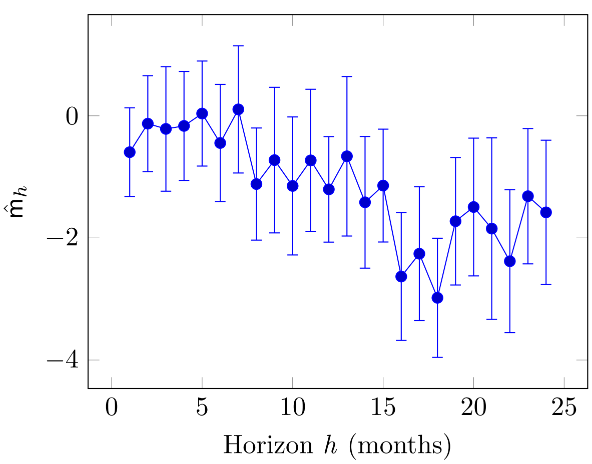

(c) Money pump (M), federal funds factor

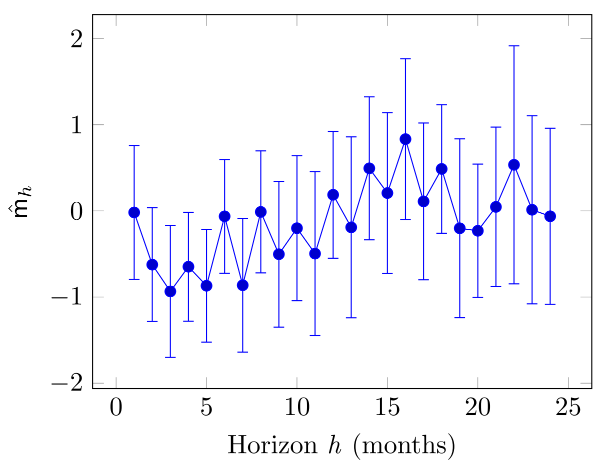

(d) Money pump (M), policy news factor

Notes: The figures show the estimated coefficients $\hat{\mu}_h$. Markers represent point estimates; whiskers indicate 95% confidence intervals using robust standard errors. The sample period underlying the estimates is monthly, 2001 to 2011. Observations for each regression range between 105 and 80 (recall 4 months are dropped from the sample because of sampling changes) sample (so a positive monetary policy shock lowers inflation). However, there appear to be two exceptions: during the early part of the sample in 2001 and during months around the financial crisis of 2007–2009 (the GFC). In this period, monetary policy, particularly changes in the federal funds rate, appears to have increased inflation and its volatility, with the effects being felt up to 24 months afterwards. Finally, federal funds rate shocks after 2008, which is part of the zero lower bound period, appear to have little effect on inflation, which suggests that there was little expectation of policy rate changes during this period. There is some evidence that policy news shocks during 2009 affected future inflation, but by 2010 there is little evidence that policy news did so.

Figure 10: Cumulative impulse responses

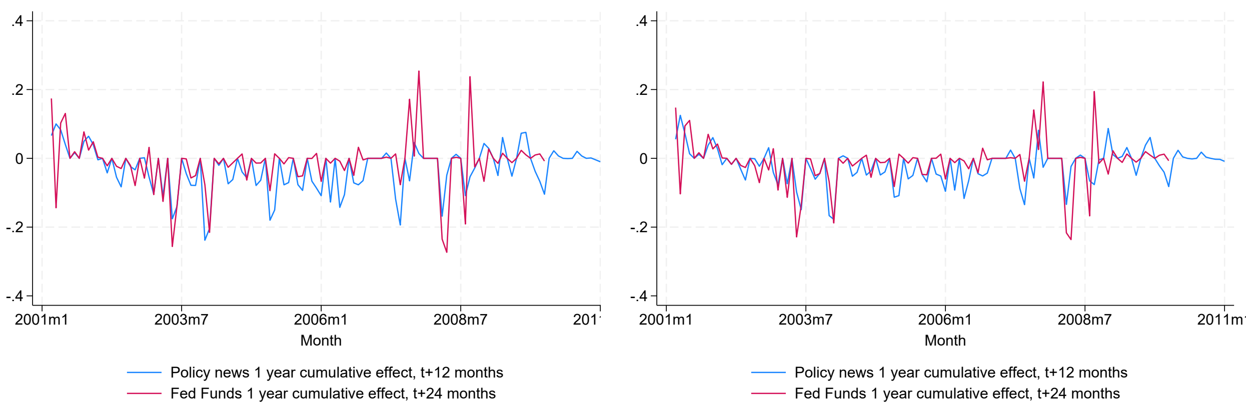

(a) Cumulative average inflation

(b) Cumulative inflation IQR

Turning to the effects of monetary policy shocks for the dispersion of inflation, the estimated coefficients for the policy news shocks and the federal funds shocks are qualitatively similar to those for average inflation. Policy news shocks primarily affect the inter-quartile range of inflation in the period 1–12 months after the shock, though the evidence for this effect is weaker both statistically and quantitatively than for average inflation. Conversely, the effect of federal funds shocks on the inter-quartile range is only marginally weaker quantitatively. Figure 10b plots the contribution of both to the inter-quartile range over the two cumulative windows considered. Similar to average inflation, monetary policy shocks tend to lower the dispersion of inflation except for the monetary policy shocks during the 2001 and 2007–2009 periods.

One policy implication of the estimated impulse response functions is that by focusing on a fixed-basket consumer price index, central banks may miss an important transmission channel, the money pump channel, for how their policy actions affect household inflation. Fixed-basket indices by construction impose $\mathsf{M}_t^i = 0$ for all i, t and so miss the contribution to household inflation from changes in money pump costs. Because these changes also offset the nominal-real impacts for both the federal funds rate shock and the policy news rate shock, then central banks are at risk of incorrectly estimating the impact of their policy actions.

A second, and perhaps more substantive, policy implication is that central bank communication appears to operate at a shorter-run horizon than interest rate decisions and also affects household inflation through two competing channels: nominal-real expenditure and preferences. Interestingly in this study, the effect of central bank communication on the behaviour channel of inflation appears to have the largest quantitative magnitude over a 1–12 month horizon. Whether the relative strength of the behaviour channel would be similar for a wider, or different, basket of consumption goods is an open question. Certainly, monetary policy inflation targets often exclude food from their expenditure baskets because food prices are thought to be

2010—that is, the cumulative effect for the period January 2009 to January 2010 caused by the shock in January 2008.

Table 6: Cumulative inflation multipliers

	1–12 months		13–24 months	
	Average	IQR	Average	IQR
$\hat{F}_t^{ER}(ffr)$	-2.046	0.764	10.76^{*}	9.291^{***}
	3.479	2.590	4.995	2.176
$\hat{F}_t^{ER}(policy)$	5.733^{**}	3.121	-0.211	-3.177
	1.673	1.743	2.889	1.889
$\hat{F}_t^{M}(ffr)$	2.248	-0.685	-11.03^{*}	-9.608^{***}
	3.250	2.393	4.850	2.133
$\hat{F}_t^{M}(policy)$	-6.817^{***}	-4.202^{*}	-0.318	2.874
	1.741	1.783	2.995	2.005
Observations	81	81	69	69

Notes: ffr refers to the federal funds rate shock; policy refers to the policy news shock. Standard errors reported in second row for each variable, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

volatile and not directly influenced by policy. Investigating whether the effects of monetary policy identified in this paper hold true for "core" consumption items would appear to be an important next step.

A third observation is that federal funds rate changes appear to affect both average household inflation and the dispersion of household inflation. There is some evidence that policy news shocks lower the dispersion of household inflation through the behaviour channel at the 1–12 month horizon but no evidence of a significant effect at the longer horizon. This suggests that central bank policy news shocks may be well communicated to all households relatively equally. However, the same would not appear to be true for federal funds rate shocks. Whether this reflects intentional decisions by some households to subject themselves to interest rate volatility or not remains an open question that cannot be addressed with the data used in the current study.

6 Conclusion

The analysis in this paper makes three contributions to understanding household inflation dynamics using a panel data of scanner purchases by households. Using results from Echenique, Lee, and Shum (2011), I show that household inflation can be decomposed into three components: nominal expenditure changes, quantity changes, and money pump costs, which reflects household behaviour. I show that household behaviour as measured by money pump costs is an important driver of household inflation. Considering only households whose preferences do not violate WARP leads to measured monthly inflation rates well above 5 percent for the non-perishable groceries for the sample examined here. Finally, I examine the effect of monetary policy on household inflation rates and show that money policy acts on household inflation through household behaviour and that central bank communications appear to affect inflation at a shorter time horizon than federal funds rate changes. The central lesson is that preferences matter for household inflation and the conduct of monetary policy.

References

Bai, J., and S. Ng (2002): "Determining the Number of Factors in Approximate Factor Models," *Econometrica*, 70(1), 191–221.

Baqaee, D., and E. Farhi (2022): "Supply and Demand in Disaggregated Keynesian Economies with an Application to the COVID-19 Crisis," *American Economic Review*, 112(5), 1397–1436.

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

Bronnenberg, B. J., M. W. Kruger, and C. F. Mela (2008): "Database paper: The IRI marketing data set," *Marketing Science*, 27(4), 745–748.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999): "Chapter 2 Monetary policy shocks: What have we learned and to what end?," vol. 1 of *Handbook of Macroeconomics*, pp. 65–148. Elsevier.

Cochrane, J. H., and M. Piazzesi (2002): "The Fed and Interest Rates - A High-Frequency Identification," *American Economic Review*, 92(2), 90–95.

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

Cook, T., and T. Hahn (1989): "The effect of changes in the federal funds rate target on market interest rates in the 1970s," *Journal of Monetary Economics*, 24(3), 331–351.

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

Echenique, F., S. Lee, and M. Shum (2011): "The Money Pump as a Measure of Revealed Preference Violations," *Journal of Political Economy*, 119(6), 1201–1223.

Eichenbaum, M. (1992): "'Interpreting the macroeconomic time series facts: The effects of monetary policy': by Christopher Sims," *European Economic Review*, 36(5), 1001–1011.

Einav, L., E. Leibtag, and A. Nevo (2008): "Not-so-Classical Measurement Errors: A Validation Study of Homescan," Working Paper 14436, National Bureau of Economic Research.

Hobijn, B., and D. Lagakos (2005): "Inflation Inequality in the United States," *Review of Income and Wealth*, 51(4), 581–606.

Hobijn, B., K. Mayer, C. Stennis, and G. Topa (2009): "Household inflation experiences in the U.S.: a comprehensive approach," Working Paper Series 2009-19, Federal Reserve Bank of San Francisco.

Jordà, Ò. (2005): "Estimation and inference of impulse responses by local projections," *American Economic Review*, 95(1), 161–182.

Kaplan, G., and S. Schulhofer-Wohl (2017): "Inflation at the household level," *Journal of Monetary Economics*, 91, 19 – 38, The Swiss National Bank/Study Center Gerzensee special issue title is "A Conference in Honor of Robert King".

Kelly, B., and S. Pruitt (2015): "The three-pass regression filter: A new approach to forecasting using many predictors," *Journal of Econometrics*, 186(2), 294–316, High Dimensional Problems in Econometrics.

Kuttner, K. N. (2001): "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market," *Journal of Monetary Economics*, 47(3), 523–544.

Lauper, C., and G. Mangiante (2021): "Monetary policy shocks and inflation inequality," Cahiers de Recherches Economiques du Département d'économie 21.02a, Université de Lausanne, Faculté des HEC, Département d'économie.

Miranda-Agrippino, S., and G. Ricco (2021): "The Transmission of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, 13(3), 74–107.

Montiel Olea, J. L., and M. Plagborg-Møller (2021): "Local projection inference is simpler and more robust than you think," *Econometrica*, 89(4), 1789–1823.

Nakamura, E., and J. Steinsson (2018): "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect," *The Quarterly Journal of Economics*, 133(3), 1283–1330.

Ramey, V. (2016): "Chapter 2 - Macroeconomic Shocks and Their Propagation," vol. 2 of *Handbook of Macroeconomics*, pp. 71–162. Elsevier.

Ramey, V. A., and S. Zubairy (2018): "Government spending multipliers in good times and in bad: evidence from US historical data," *Journal of Political Economy*, 126(2), 850–901.

Redding, S. J., and D. E. Weinstein (2019): "Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences*," *The Quarterly Journal of Economics*, 135(1), 503–560.

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

Sims, C. A. (1992): "Interpreting the macroeconomic time series facts: The effects of monetary policy," *European Economic Review*, 36(5), 975–1000.

Smeulders, B., L. Cherchye, F. C. R. Spieksma, and B. D. Rock (2013): "The Money Pump as a Measure of Revealed Preference Violations: A Comment," *Journal of Political Economy*, 121(6), 1248–1258.

Stock, J. H., and M. W. Watson (2002): "Forecasting using principal components from a large number of predictors," *Journal of the American Statistical Association*, 97(460), 1167 – 1179.

——— (2012): "Generalized Shrinkage Methods for Forecasting Using Many Predictors," *Journal of Business & Economic Statistics*, 30(4), 481–493.

7 Appendix

7.1 Numerical example

It may be instructive to illustrate the role of money pump costs for inflation with a simple example. Suppose $\mathbf{x}$ comprises apples and oranges. In period 1, apples cost \$2 and oranges cost \$1. Given these prices, household j purchases 3 apples and 2 oranges for a cost of \$8. In period 2, prices change to \$1 and \$2 respectively, and the same household purchases 1 apple and 4 oranges for a cost of \$9. This clearly violates WARP since the household revealed a preference for apples relative to oranges in period 1 and yet purchases fewer apples in period 2 even though the relative price falls. Using the definition of M_t^j above, it is straightforward to calculate that $\mathsf{M}_t^j = 4$ in this example:[17]

$$M_t^j = \begin{bmatrix} 2 & 1 \end{bmatrix} \begin{bmatrix} 3-1 \\ 2-4 \end{bmatrix} + \begin{bmatrix} 1 & 2 \end{bmatrix} \begin{bmatrix} 1-3 \\ 4-2 \end{bmatrix} = 2 + 2 = 4.$$

If all households in the economy are households of type j, then money pump costs deflated by first-period expenditure are $\frac{4}{8}$ and average (or aggregate) household inflation is, using the decomposition from Equation (2):

[17] One interpretation of M_t^j proposed by Echenique, Lee, and Shum (2011) is that it is the amount of money a fictitious and devious arbitrageur could earn from exploiting the household's purchases. If the arbitrageur were able to purchase an apple from period 2 at \$1 to sell in period 1 and buy oranges at \$1 to sell in period 2, the arbitrageur would earn \$4.

$$\pi_t^i = \; -\frac{1}{8} \; = \; \frac{9}{8} \; - \; \frac{6}{8} \; - \; \frac{4}{8},$$

where the nominal expenditure growth is $\frac{9}{8}$ and real expenditure growth is $\frac{6}{8}$. In this simple example, the money pump cost contributes -0.5 points to the household's overall inflation rate of -0.125, and the households experience deflation despite nominal expenditure growth being higher than real expenditure growth.

In a different economy populated by households of type k that instead purchased the reverse bundles at the given prices, $M_t^k = -4$ and households would have spent \$6 and \$7 in periods 1 and 2, respectively (and not violated WARP). These households would have experienced an inflation rate of:

$$\pi_t^i = \; \frac{1}{2} \; = \; \frac{7}{6} \; - \; \frac{8}{6} \; - \; \frac{-4}{6},$$

despite real expenditure growth being higher than nominal growth.

It is also straightforward to show arithmetically that the composition of an economy may matter for inflation and aggregate expenditure. For example, an economy populated by equal numbers of households of type j and k would have an aggregate money pump cost of:

$$\frac{\mathsf{M}_t^j}{\mathbf{p}_{t-1}^i \mathbf{x}_{t-1}^j} + \frac{\mathsf{M}_t^k}{\mathbf{p}_{t-1}^i \mathbf{x}_{t-1}^k} = \frac{4}{8} + \frac{-4}{6} = -\frac{1}{6} \neq 0.$$

In an economy with equal numbers of j and k households, then, average inflation would be higher because money pump costs are negative.

In this stylized example, the population proportion of households j with $M = 4$ and k with $M = -4$ must be 4/7 and 3/7, respectively, for money pump costs to have no effect on average inflation. However, simply eliminating M is not, by itself, sufficient for average household inflation to track the fixed-basket measure because of differences in household consumption bundles. Without loss of generality, assume that the fixed-basket measure of inflation is 0. For $\frac{1}{N}\sum_{i=1}^{N} \frac{(\mathbf{p}_t^i - \mathbf{p}_{t-1}^i)\mathbf{x}_t^i - \mathsf{M}_t^i}{\mathbf{p}_{t-1}^i \mathbf{x}_{t-1}^i} = 0$, it requires population proportions of 20/41 and 21/41, respectively, for types j and k.

This simple example generates a relatively high value of M_t^j for household i in proportion to the base-period expenditure, $4/8 = 50\%$ of the purchase value spent by the household in $t-1$, which implies that the last term in Equation (2) has a large effect on household inflation. Whether empirical values of M_t^i are sufficiently large to imply much difference in the time series of inflation appears to be an open question.

7.2 Charts

Figure 11: Using alternative measures of counterfactual prices

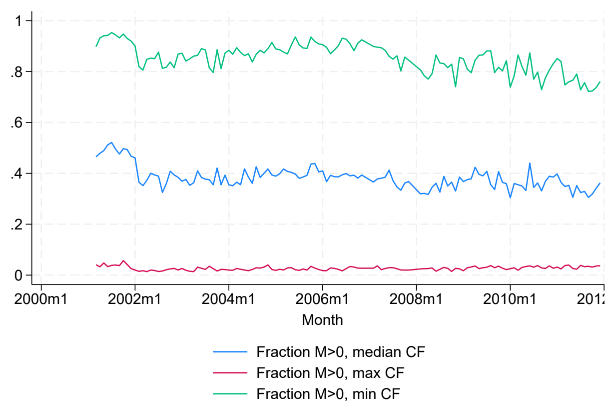

(a) Fraction $M > 0$ by counterfactual price assumption

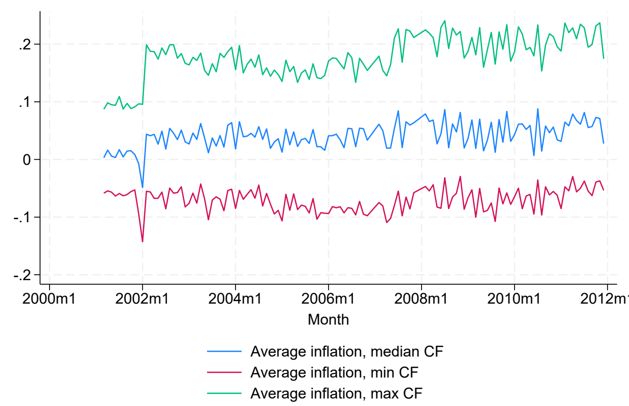

(b) Household inflation by counterfactual price assumption

Figure 12: Monthly average household inflation and IQR

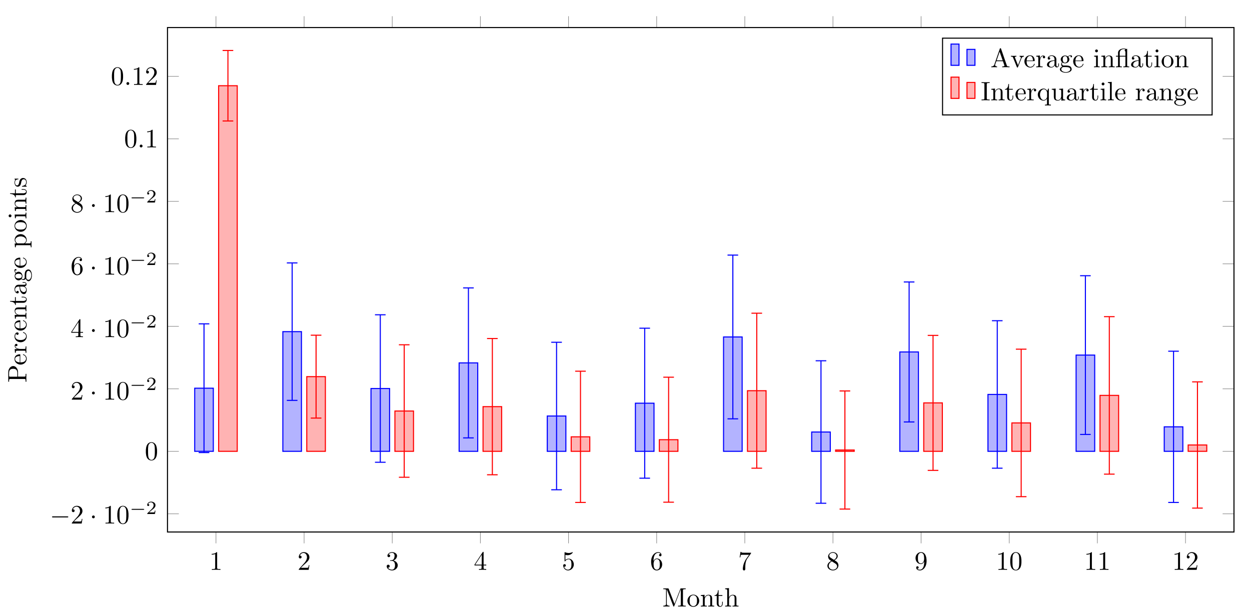

Notes: The figures show the estimated coefficients from monthly dummy variables for $\bar{\pi}_t$ and $\iota(\pi)_t$ over the sample period from February 2001 to December 2011. Note that the first month, January, also includes the average over the full sample.