Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

A New Reason to Hate Grocery Inflation: Measuring and Interpreting Inflation Heterogeneity

Kelsey O'Flaherty

2026-001

Please cite this paper as:
O'Flaherty, Kelsey (2026). "A New Reason to Hate Grocery Inflation: Measuring and Interpreting Inflation Heterogeneity," Finance and Economics Discussion Series 2026-001. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2026.001.

A New Reason to Hate Grocery Inflation: Measuring and Interpreting Inflation Heterogeneity*

Kelsey O'Flaherty

Board of Governors of the Federal Reserve System
20th Street & Constitution Avenue NW
Washington, DC 20551

Phone: 240-975-0972
E-mail: kelsey.j.oflaherty@frb.gov

Abstract

The 2021-2022 inflation episode presented the first opportunity to examine inflation and price dispersion using U.S. scanner data in a high-inflation environment. Data from 50,000 outlets reveals that price changes across similar goods grew more dispersed in 2022 before falling again in 2023. This paper documents how price change dispersion interacts with households' product choices to generate substantial inflation heterogeneity. Household-level inflation rates exhibit a 1.4 percentage point interquartile range in 2019, which grew to 4.0 percentage points in 2022 before falling back to 1.6 percentage points in 2023. Households offset little of their implied budget shocks through substitution. A model with idiosyncratic preferences rationalizes household behavior and implies that households' inflation rates represent convenient, observable bounds on their welfare losses. When inflation peaked in 2022, households at the 10th and 90th percentiles of the inflation distribution and average grocery expenditures faced welfare losses of $573 and $1,145, respectively.
JEL Codes: E31, E21, D11, D12

*__Acknowledgements__: I gratefully acknowledge financial support of the Re-Engineering Statistics using Economic Transactions (RESET) project, a joint project of the University of Maryland, the University of Michigan, and the U.S. Census Bureau. RESET is supported by the Alfred P. Sloan Foundation (G-2016-7227, G-2021-16821). This project also benefited greatly from the feedback of Matthew Shapiro, Joshua Hausman, John Leahy, and Melvin Stephens; from discussions with Aditya Aladangady, Costas Arkolakis, Javier Cravino, Gabriel Ehrlich, Sun Kyoung Lee, Jacob Orchard, Francisca Sara-Zaror, and others, including over the course of presentations on the job market.

Many households report pain from inflation which far exceeds a standard model's prediction based on temporary gaps between average price increases and wage gains.[1] This paper provides new evidence that households face dramatically different inflation rates, some consistent with large welfare losses, and shows why a typical framework for evaluating the costs of inflation fails to capture these experiences. Much of the literature documenting inflation heterogeneity focuses on differences in shopping behavior, location, or the relative importance of each category of goods in households' consumption baskets (for example, prescription drugs for older households). Those factors tend to generate differences in annualized inflation of less than 0.5 percentage points across income or other demographic groups. I find that price change heterogeneity across individual items within product categories gives rise to much larger differences in inflation across households. Households' idiosyncratic inflation experiences arising from this mechanism grew more disparate in 2022, with 10 percent of households experiencing grocery inflation almost 4 percentage points above average. It perhaps should not surprise if such households report hating the recent episode of grocery inflation: it caused them significantly more financial distress than aggregate or group inflation statistics imply.

To track the price changes of individual goods and how household purchasing patterns respond to those changes, I use two rich datasets from NielsenIQ. The first, its Consumer Panel, tracks the grocery purchases of about 60,000 households using in-home barcode scanners to record the quantity and price paid for each item purchased. A second, the NielsenIQ Retail Scanner data, contains weekly summaries of the revenues and units sold for each barcode at about 50,000 participating outlets. The spending categories well-captured by this dataset (those entering my "grocery" definition) cover approximately 12% of expenditures.[2] The level of granularity in this data, its extensive coverage of the grocery sector, and its large panel of households are all essential for my empirical strategy.

In the first examination of price change dispersion using U.S. scanner data during a period of high inflation, I find that the interquartile range of price changes among individual items within detailed product categories rose for almost every product type between 2019 and 2022 before receding somewhat in 2023. As an illustrative example of this finding, the interquartile range of four-quarter price changes across individual coffee varieties—for instance Peet's French Roast or Folger's Classic Roast—rose from 2 percentage points in 2019 to 15 percentage points in 2022. In the latter period, one quarter of households saw the price of the coffee variety they purchased in 2021 rise by 5 percent or less while another quarter saw the price of their coffee rise by over 20 percent. The HomeScan consumer panel captures how households respond to such widely varying price changes. Examining substitution patterns at the household level through 2023, I find that households whose chosen varieties undergo large price increases do not tend to meaningfully mitigate these shocks by switching to goods with smaller price changes.

These findings point to the individual product varieties a household purchases being a major

[1] See "The Elusive Costs of Inflation..." (Nakamura et al., 2018) for a survey of the costs of inflation and Georgarakos et al. (2025) for a discussion of households' direct reports of the same.

[2] Table A.1 in the Supplemental Appendix identifies these categories.

determinant of its inflation rate. Most inflation analyses ignore households' individual product choices, either by explicitly focusing on an aggregate bundle of goods or by assuming a common inflation rate among homogeneous goods like coffee. Those assumptions would be innocuous if similar items exhibited similar price changes or if households faced minimal welfare losses switching from products with above-average inflation to substitutes with slower price growth. Having found neither assumption to be true, my paper suggests a new characterization of how households experience inflation.

To address the measurement implications of these findings, I provide new estimates of household-level inflation which account for this reality that households purchase distinct product varieties which may experience differing price changes. These measures overcome a significant missing data problem that arises when using barcode-level data to compute household inflation. At this level of detail, household consumption baskets exhibit a large degree of product churn. Products enter and exit the basket as households switch brands, change flavors, or when manufacturers change product attributes like size or labeling (which induces a barcode change). Because a household's price change is unobserved for goods that enter or exit its basket, previous measures of household inflation exclude such items and consider only matched goods. Unfortunately, those goods cover only about one third of grocery expenditures. More problematically, they by definition exclude goods households substitute toward or from, making it difficult to evaluate whether households offset high idiosyncratic inflation through substitution. I instead address this missing data problem by inferring the price changes households faced for goods that enter or exit its consumption basket. The NielsenIQ data often allow such price changes to be inferred at the specific outlet or retailer where the household shops, or else in its local market area.

Baseline household inflation measures in this paper provide a conservative underestimate of inflation heterogeneity while minimizing the influence of missing data. They do this by assuming that all households face the national average price for each barcode. This approach isolates the role of households' product choices from additional, well-documented heterogeneity in the prices households pay for a given barcode due to geography or shopping behavior. Though these baseline estimates aim to provide a lower bound on the extent of heterogeneity across households, the estimates reflect large disparities. In 2019, the distribution of households' Q4/Q4 grocery inflation spanned an interquartile range of 1.4 percentage points. Grocery inflation became more dispersed after 2021, with this interquartile range more than doubling in 2022 to 4 percentage points. Even in the low-inflation environment of 2019, such cross-sectional differences are large relative to both time-series variation in grocery inflation and relative to estimates of how inflation varies across income and other demographic groups. The latter discrepancy reflects aggregation bias: constructing the representative bundle for a demographic group averages across many items, masking the heterogeneity that arises from idiosyncratic exposure to specific products.

In a final contribution, the paper proposes a framework to rationalize and interpret household-level inflation rates. In the spirit of McFadden (1974), the model infers that households not substituting away from items with large price changes have strong idiosyncratic preferences for those goods. The model connects the strength of those preferences to a CES representative consumer, with empirically relevant aggregate elasticities of substitution generating significant product commitments for some consumers. I demonstrate that households' idiosyncratic inflation arising from such product commitments have welfare consequences, with household-level price indices providing non-parametric, dollar-denominated bounds on how much more households have to pay following a price change to reach the same utility brought by their original grocery bundle. These budget shocks are large and heterogeneous: in the low-inflation period of 2019, a household with average grocery spending and a 90th percentile inflation rate would need to spend approximately \$170 more than they had one year earlier to achieve the same satisfaction, while a household experiencing 10th percentile inflation would be able to do so with a basket costing about \$25 *less* than in the previous year. That gap of \$195 grew to \$572 by the end of 2022. Basket price shocks exhibit limited mean reversion, implying that households' budget shocks are moderately persistent. In a simple intertemporal consumption model, a 90th percentile grocery inflation draw in 2022 implies lifetime utility losses 0.2 percent higher than for households at the 10th percentile.

Section I of the paper begins by discussing what is already known about price dispersion and inflation heterogeneity, and it highlights questions which have not yet been addressed. Section II provides detail on the NielsenIQ HomeScan and Retail Scanner data used to measure households' grocery inflation. Section III discusses the construction of household inflation measures and what they imply about household substitution patterns. Section IV presents the key empirical results of the paper, and Section V describes the model used to interpret these results. Section VI assesses the costs of idiosyncratic inflation in that framework and documents how households respond to high idiosyncratic inflation rates. Consistent with the model's implication that household inflation rates reflect genuine budget shocks, households respond to their personal inflation by reducing consumption to offset approximately one-half of the implied budget shock. Section VII concludes and points to areas for future research.

I. Literature

A large literature examines how inflation and the cost of living differs across households. This section begins by briefly discussing the paper's connection to two strands of that literature: average inflation differences by demographic group and heterogeneity in prices paid for similar items. It then describes its empirical contribution to a third strand: inflation heterogeneity across individual households. I also note important connections to recent contributions in the literature on macroeconomic costs of inflation.

Until recently, work on inflation heterogeneity focused on average differences across demographic groups such as age, income, or family size. Such an approach begins by constructing a representative basket of goods for households in the given group. That basket might reflect the greater importance of gasoline for low-income households or medication for older households. Because cost growth or price volatility is larger in some sectors than others, this basket effect can generate inflation disparities. Jaravel (2024) provides an important recent example of this approach, using publicly available Consumer Expenditure Survey and Consumer Price Index data to construct price indices by income, age, race, and area population. Between 2002 and 2019, inflation rates of the top- and bottom-quartile income groups differed by approximately 0.25 percentage points annually. A slightly smaller gap of 0.22 percentage points annually separates the 65-74 year age group from those aged 25-34, though those age 75 and over experience considerably higher inflation. Differences were negligible across racial groups or rural/urban households.[3] In a study also using NielsenIQ data, Argente and Lee (2021) find differences of a similar magnitude in grocery inflation. Households with income above $100,000 had inflation rates 0.33 percentage points lower on average than households with income below $25,000 from 2004-2016.

In examining household-level inflation rates, this paper most closely relates to Kaplan and Schulhofer-Wohl (2017). These authors were the first to use a detailed consumer panel to construct household grocery baskets, and their results were striking. From 2004 to 2013, inflation rates have an interquartile range of between 6 and 9 percentage points across households, making the aggregate inflation rate "almost irrelevant as a source of variation in the household-level inflation rate" (p. 35). Demographics explain almost none of this variation, and two thirds of heterogeneity arises due to households paying different prices for the same barcodes. A series of concurrent papers apply their methodology in the recent inflation episode, corroborating these core points. Chen et al. (2024) provide an excellent analysis of the United Kingdom; Kostyshyna and Ouellet (2024) and Pathak-Chalise (2025) examine Canada and the United States.

The use of detailed consumption data to compute household-level inflation creates an important methodological challenge that has plagued subsequent work. Given the centrality of household-level price heterogeneity in their findings, Kaplan and Schulhofer-Wohl consider only the set of "matched" goods for which they can directly compute households' price changes: goods purchased in both the the base and reference periods. That methodology inhibits interpretation in two ways. Kaplan and Schulhofer-Wohl report that matched goods account for only about 20% of household spending in their sample, making it difficult to distinguish heterogeneity from the noise of a small sample.[4] More importantly, focusing only on matched goods also makes it impossible to quantify the extent to which households react to inflation by substituting toward cheaper alternatives as relative prices change. By construction, such observations are omitted from analysis.[5]

This paper constructs conservative estimates of inflation heterogeneity which incorporate over 90% of expenditures in the categories considered, including those that households substitute from or toward.[6] Its baseline estimates accomplish this by assuming that all households encounter the national average price change for each barcode. Though extreme in removing a genuine source of

[3] Each of these differences almost exactly correspond to the results of McGranahan and Paulson (2005) who run the same exercise from 1982-2004.

[4] Note, if this low number surprises, that different flavors and pack sizes constitute different goods.

[5] This reality limits the interpretability of the substitution analysis presented in Pathak-Chalise (2025, §6) and Kostyshyna and Ouellet (2024, §3.2). As in this paper, Chen et al. (2024) apply average prices in their baseline results (though they do not discuss the methodological significance of that choice).

[6] The introduction and phase-out of goods implies some portion are not transacted, and therefore don't have an observable price, in one of the quarters being compared. Goods carried only by outlets not included in the NielsenIQ Retail Scanner data are also excluded if not purchased by a sufficient number of households in the Consumer Panel.

heterogeneity, such an approach allows this paper to clarify that large inflation disparities persist even when considering nearly the full grocery basket, and that the opportunity to substitute away from high-priced goods does little to mitigate those disparities.

Lastly, the paper adds to an active literature on the economic costs of inflation in two ways. It provides new evidence from U.S. scanner data that price changes became more dispersed during the recent period of high inflation, and it formalizes a channel by which aggregate inflation affects households: one based on the interaction between average inflation and cross-sectional budget shocks. In many macro models, the economic cost of inflation depends on the degree of price dispersion inflation generates. If inflation generates more inefficient price dispersion, consistent with a Calvo pricing mechanism, inflation has much larger economic costs than if prices update more flexibly. Work by Nakamura et al. (2018) and Alvarez et al. (2019) empirically falsified a key prediction of the Calvo pricing mechanism—that the absolute size of incremental price changes become larger in periods of high inflation—implying prices should be flexible enough to make inflation costs low for reasonable inflation rates. In a direct examination of price level dispersion using web-scraped prices, Cavallo and Kryvtsov (2024) document a compression of prices as low-priced goods experienced higher inflation. They note, however, that the welfare implications of this decline in price dispersion may deviate from the model-based framework as households trading down the quality ladder face higher prices (p. 2-3, 10).

My paper formalizes this cost of inflation, presenting a framework that relates the budget shocks households endure from a given inflation to the dispersion of price *changes* they encounter. Downstream of the price setting process, a wider distribution of inflation outcomes implies that a larger share of households face real increases in the cost of living. If the relationship observed between inflation and its cross-sectional variance from 2021 through 2023 reflected a consistent pattern, even a "neutral" aggregate inflation perfectly offset by higher average wage gains would generate welfare losses.[7] This paper applies standard household utility assumptions to demonstrate that those costs rise sharply with the spread of household inflation experiences and may have been non-trivial in the recent inflation episode.

II. Data

NielsenIQ provides two datasets ideal for computing household grocery inflation. Its Consumer Panel tracks the purchases of approximately 60,000 households over the period considered. That dataset captures the quantity and price paid for each item, providing a detailed account of consumption baskets at the barcode level. The NielsenIQ Retail Scanner data complements the Consumer Panel with weekly summaries of the revenues and units sold for each barcode at participating outlets covering about half of the market for grocery and drug stores, as well as about one third of the market for mass merchandisers. Table 1 summarizes the number of households and outlets captured by these datasets and used in this paper's analysis sample.

[7] Orchard (2020) raises this point as well, having documented a similar co-movement between aggregate inflation and household-level inflation using sectoral price changes through 2017.

NielsenIQ selects its Consumer Panel households to be representative of the population, providing projection factors to account for differential participation. These households record their purchases from each shopping trip using in-home barcode scanners. After identifying the store, households record the quantity of each item purchased. If the trip was to a store covered by the Retail Scanner data, the household does not enter a price. Instead, NielsenIQ automatically applies the average price paid at its store during the week of purchase. Otherwise, households enter the total expenditures for each good they scan. All households further record the value of coupons applied to each good and whether they perceived the item to be on sale.[8] The sample of households used in this paper come from a subset of panelists asked to record Magnet data, which includes products without barcodes such as fresh vegetables. For inclusion in any quarter's inflation analysis, households must also record at least $150 in included expenditure categories during the given quarter, the quarter preceding it, and four quarters prior.[9] Households may continue to participate in the Consumer Panel for any duration, but about 20% attrite annually.

Because the Consumer Panel data records the outlet of each purchase, and many of these same outlets appear in the Retail Scanner data, a significant portion of household expenditures can be linked to administrative data. Linking the datasets is central to this paper's contribution. It offers a way to overcome the matched goods problem of Kaplan and Schulhofer-Wohl by inferring price changes for the goods a household does not purchase in both periods. Table 2 summarizes the amount of spending households record in the Consumer Panel and how much of that spending meets various criteria. Matched goods account for a bit more than a quarter of recorded expenditures. By combining NielsenIQ HomeScan and Retail Scanner data, however, about two-thirds of recorded expenditures can be linked to the retailer where the purchase was made.

Products in the NielsenIQ data are categorized into groups, the most detailed of which (called product modules) represent approximately 1,100 homogeneous items such as frozen chicken, ground coffee, or men's shaving cream. These goods are further organized into product groups and departments.[10] Analysis in this paper uses spending data in the dry grocery, frozen food, dairy, deli, packaged meat, fresh produce, alcohol, health/beauty care, and other non-food grocery departments capturing products like cleaning supplies. Not included are purchases in the general merchandise department, which includes items like electronics, small household appliances, and cookware. Excluding these categories prevents large and infrequent purchases from affecting estimates of inflation heterogeneity. The remaining categories cover approximately 80% of recorded expenditures. Among those categories, price changes can be computed at some level for 92% of expenditures. The core results of this paper rely on national average price changes for each barcode. Supplemental Appendix A.1 discusses the computation of these price changes, an endeavor which

[8] Automatic price reporting at NielsenIQ retailers necessarily mixes the prices paid by households applying or not applying coupons or taking advantage of customer rewards programs.

[9] Table A.1 of the Supplemental Appendix enumerates included products, and Section A.1 discusses additional sample selection details.

[10] In 2021, NielsenIQ overhauled its classification system such that goods no longer fall into mutually exclusive categories. Through at least 2023, however, mappings make it possible to continue classifying UPCs into their original product groups and modules. For the sake of consistency, I apply the pre-2021 classification scheme for all analysis.

aims to optimally leverage the representativity of the HomeScan purchases and the broad coverage of items in the Retail Scanner data.

III. Measuring Household Inflation

Although aggregate price indices can indicate broad price trends, the theoretical assumptions necessary to assign them welfare relevance to all consumers are implausible.[11] Further, the concept price indices ought to measure—changes in the cost of achieving a fixed standard of living—is not necessarily well-defined for large groups of heterogeneous individuals. The definition given by Pollak (1989) is illustratively specific. Define the price index $\Psi_{h,t-j,t}$ for individual h between periods $t-j$ and t as the ratio of spending $S()$ needed under price vectors $\vec{p}_t$ and $\vec{p}_{t-j}$ to achieve base utility U_{h0} given base preferences R_{h0}

$$\Psi_{h,t-j,t} = \frac{S(\vec{p}_t, R_{h0}, U_{h0})}{S(\vec{p}_{t-j}, R_{h0}, U_{h0})}. \tag{1}$$

Attention to subscripts highlights the complexities making computation of this theoretical index impractical. Preferences R surely vary across individuals, and may even change over time for a given individual. Subscript 0 denotes the "base," which might (but need not) refer for instance to the preferences held in period t or $t-j$. Even if preferences were identical across households, the relationship between expenditures and satisfaction could vary with the level of utility U_{h0}.

While equation 1 illustrates the theory embedded in a price index, it offers little guidance on computing one. Exact price indices choose particular preference representations and derive formulas for how an agent's particular bundle of goods and prices paid relate to utility. In contrast, statistical price indices rely on consumer behavior assumptions to infer how each good's price change affects household well-being.

This paper takes the latter approach, computing Tornqvist price indices for each household in its headline results. The Tornqvist index is a geometric average of two other price indices which I later demonstrate act as bounds on the theoretical change in the cost of living under mild behavioral assumptions. As an average of these bounds, the Tornqvist falls into a class of price indices characterized as "superlative." Such indices are exact for flexible utility functions which can second-order approximate any twice-differentiable, linearly homogeneous function (Diewert, 1976, p. 117). Put simply: whatever the preferences of household h may be, Equation 2 is typically assumed to represent a reasonable approximation for its change in the cost of living in the region of expenditures being compared.

The geometric Tornqvist index is defined as

$$\Psi^T_{h,t-j,t} = \sqrt{\Psi^L_{h,t-j,t}\Psi^P_{h,t-j,t}} \tag{2}$$

[11] Deaton and Muellbauer (1980, §7.1, p.170-178) describe the theory summarized here and discuss the core assumptions needed for a representative consumer to have a price index with welfare relevance to all households.

where $\Psi^L_{h,t-j,t}$ and $\Psi^P_{h,t-j,t}$ are the geometric Laspeyres and Paasche indices, respectively. A Laspeyres index weights the importance of each good k according to its expenditure share in the $t-j$ consumption basket of household h. Denote the set of all goods k in that basket for which a relative price can be defined as $\Omega_{h,t-j}$.[12] The equation of the geometric Laspeyres index is

$$\text{(3)} \qquad \Psi^L_{h,t-j,t} = \prod_{k \in \Omega_{h,t-j}} \left(\frac{p_{h,k,t}}{p_{h,k,t-j}} \right)^{s_{h,k,t-j}} \quad \text{where} \quad s_{h,k,t-j} = \frac{p_{k,h,t-j} q_{k,h,t-j}}{\sum_{\Omega_{h,t-j}} p_{k,h,t-j} q_{k,h,t-j}}.$$

A Paasche index parallels this construction, but weights each good by its share in t:

$$\text{(4)} \qquad \Psi^P_{h,t-j,t} = \prod_{k \in \Omega_{h,t}} \left(\frac{p_{h,k,t}}{p_{h,k,t-j}} \right)^{s_{h,k,t}} \quad \text{where} \quad s_{h,k,t} = \frac{p_{k,h,t} q_{k,h,t}}{\sum_{\Omega_{h,t}} p_{k,h,t} q_{k,h,t}}.$$

In practice, inflation rates tend to be estimated using national expenditure shares $\bar{s}_{k,t-j}$ or $\bar{s}_{k,t}$ and changes in national aggregate prices $\bar{p}_{k,t}$ (or urban averages in the case of the CPI). For statistical agencies, this reflects practical realities of data collection costs and a priority for releasing timely estimates that capture national trends. In principle these statistics have no more relevance to a given household's cost of living than GDP estimates have for that household's income.

To estimate price indices at the household level, Hobijn and Lagakos (2005) and McGranahan and Paulson (2005), as in most of the early work on price indices, use Consumer Expenditure Survey microdata to identify household-specific expenditure shares $s_{h,k,t-j}$. Their approaches employ relatively coarse aggregate definitions of products k (e.g. "milk" in place of "Silk unsweet vanilla almond milk") and average price changes for those aggregates, making it difficult to capture substitution patterns or movements along the quality ladder within each category. Kaplan and Schulhofer-Wohl (2017) therefore introduced a major innovation by using detailed microdata in the NielsenIQ Consumer Panel data to track products at the barcode level. They also employ prices actually paid by each household, denoted $p_{h,k,t}$. A Tornqvist index so defined would be given by

$$\text{(5)} \qquad \Psi^{KSW}_{h,t-j,t} = \prod_{k \in \Omega^*_{h,t-j,t}} \left(\frac{p_{h,k,t}}{p_{h,k,t-j}} \right)^{\left(\frac{s^*_{h,k,t-j} + s^*_{h,k,t}}{2} \right)}$$

where $\Omega^*_{h,t-j,t}$ is the limited set of matched goods which household h purchases both in period $t-j$ and in t. Likewise, the shares for each good $s^*_{h,k,t}$ are computed only among this set of continuing goods. Kaplan and Schulhofer-Wohl make this choice because of an inability to define households' price changes between $p_{h,k,t-j}$ and $p_{h,k,t}$ for goods not purchased in both periods. One only knows

[12] Some price changes cannot be defined between $t-j$ and t because the good is not available due to stockouts or product turnover. In other cases, items sold only at retailers not participating in the Retail Scanner Data may not be purchased by a member of the HomeScan panel. To ease notation, I omit this dependence between $t-j$ and t from the subscripts of Ω. Goods purchased by household h in $t-j$ for which price changes can be defined between $t-j$ and t are denoted simply $\Omega_{h,t-j}$ while those bought in period t for which price changes can be defined are denoted $\Omega_{h,t}$. Likewise, it is everywhere assumed that shares $s_{h,k,t-j}$ and $s_{h,k,t}$ are computed only over these sets of good such that they sum to 1 over the sets $\Omega_{h,t-j}$ and $\Omega_{h,t}$, respectively.

what price household h paid for good k in periods when it purchases good k.

The headline measure of this paper aims to capture a conservative underestimate of inflation heterogeneity: one that reflects the full extent of basket differentiation without adding noise. It therefore applies a hybrid approach, computing Laspeyres and Paasche indices using detailed household expenditure shares but national average price changes for each item

$$\Psi^L_{h,t-j,t} = \prod_{k \in \Omega_{h,t-j}} \left(\frac{\bar{p}_{k,t}}{\bar{p}_{k,t-j}} \right)^{s_{h,k,t-j}}, \tag{6}$$

$$\Psi^P_{h,t-j,t} = \prod_{k \in \Omega_{h,t}} \left(\frac{\bar{p}_{k,t}}{\bar{p}_{k,t-j}} \right)^{s_{h,k,t}}, \tag{7}$$

and the Tornqvist average of the two

$$\Psi_{h,t-j,t} = \sqrt{\Psi^L_{h,t-j,t} \Psi^P_{h,t-j,t}}. \tag{8}$$

It will often be more transparent in what follows to consider this index in log form

$$ln\left(\Psi_{h,t-j,t}\right) = \sum_{k \in \left\{\Omega_{h,t-j} \cup \Omega_{h,t}\right\}} \left(\frac{s_{h,k,t-j} + s_{h,k,t}}{2} \right) \Delta ln(\bar{p}_{k,t-j,t}). \tag{9}$$

where the set $\{\Omega_{h,t-j} \cup \Omega_{h,t}\}$ reflects that the index incorporates all goods purchased in t or $t-j$.

Using the NielsenIQ HomeScan panel's detailed household expenditures allows these indices to capture important dynamics of item-level heterogeneity within product categories. Applying a law of one price assumption allows price indices to be inferred for all items except those that exit or enter the market nationally or go unsold in this data. That reduces the variability arising from the small sample of matched goods and permits analysis of the substitution dynamics by which households might lower their inflation rate.

Concretely, suppose that household h purchases Almond Breeze almond milk each week in 2018Q4. At some point in 2019, its price increased relative to other brands. Household h has purchased Berkeley Farms 2% milk ever since. The NielsenIQ Consumer Panel records the price of Almond Breeze each time it was purchased, which can be aggregated to an expenditure-weighted average price $p_{h,A,18Q4}$. In 2019Q4 the same quarterly average price paid by household h can be constructed for $p_{h,B,19Q4}$. A matched goods methodology removes both goods from the basket used to compute inflation. The alternative that this paper proposes instead incorporates both goods in a way that captures such a substitution. Almond Breeze appears in the Laspeyres index of equation 6 with its expenditure share in 2018Q4 $s_{h,A,18Q4}$ while Berkeley Farms 2% milk appears in the Paasche index of equation 7 with its 19Q4 expenditure share $s_{h,B,19Q4}$.[13]

[13] As Section V discusses, the Laspeyres index serves as an upper bound to the cost of this price change by assuming that the household gained no utility by substituting (or was exactly indifferent between buying Berkeley Farms milk

The resulting index uses almost the entire grocery basket and removes the possibility of volatile temporary price movements driving inflation heterogeneity. Section B of the Supplemental Appendix presents the construction of more realistic measures that do incorporate price heterogeneity.

IV. Inflation Heterogeneity

Having reviewed the fundamentals of how individual price changes contribute to households' inflation experiences, and how different indices capture households' substitution responses, the tools are in hand to evaluate key results of the paper. This section begins by documenting how price change dispersion evolved from 2021 through 2023 before examining how that phenomenon shapes households' inflation experiences. It also discusses why the role of item-level price changes have not been apparent in previous studies.

A. The Role of Detailed Price Change Dispersion

NielsenIQ Retail Scanner data provide an ideal environment to test whether item-level price changes became more dispersed when inflation rose in 2021 and 2022. In capturing the sales of every barcode at participating retailers, it provides a complete picture of how price trends of individual products evolved at major retailers over this period.

The interquartile range of four-quarter, national average log price changes within each NielsenIQ product module provides a convenient measure of dispersion with immediate relevance to the range of household inflation experiences. A simple version of that metric for a product module with 10,000 available products would rank each by the size of its price change, then difference the price changes of the 7,500th and 2,500th ranked products. Because a few products tend to dominate sales within their product categories, however, such a measure assigns outsize importance to products of little relevance to most households. Instead of assigning each product an equal weight when finding the 75th and 25th percentiles, I assign each product its share of aggregate expenditures in its category. Comparing the 75th and 25th percentiles of that distribution provides a sense of price change dispersion consumers actually encounter.

Each of the approximately 1,000 product modules used to classify sales at NielsenIQ retailers exhibits a different interquartile range of price changes among its products. Rather than reduce these individual estimates into a single statistic, each line of Figure 1 plots the spectrum of interquartile ranges across product modules for the specified year.[14] It shows that in 2019 most product modules have interquartile ranges of about 5 percentage points. In a small number of modules, the 75th percentile price change is at least 10 percentage points higher than the 25th percentile price change. Price changes became much more dispersed in 2022, with interquartile

in 19Q4 or paying the new, higher price for Almond Breeze). Conversely, the Paasche index constitutes a lower bound under the assumption that substitution was costless (the household was exactly indifferent between Almond Breeze and Berkeley Farms milk in 18Q4, before the price change).

[14] To improve visibility, it plots only modules in the 10th - 90th percentiles of price change dispersion.

ranges nearly doubling since 2019 across the distribution of product modules.[15,16] Both as inflation was rising in 2021 and as it was falling in 2023, the degree of dispersion was elevated relative to 2019 but well below the levels seen in 2022.

B. Household Substitution Responses

The wide distribution of price changes documented in Figure 1 need not imply an equally wide distribution of within-category inflation rates. Households purchasing a variety of coffee in 2021 which undergoes a 20% price change through 2022 have the opportunity to avoid that price change by substituting to a different variety, perhaps one exhibiting the average price change of 12%. Section III described how to evaluate the extent to which households reduce inflation through substitution: by checking whether households' Paasche indices are materially lower than their Laspeyres indices. Because the Paasche index evaluates the price change of the good purchased in 2022, it captures any benefit of such substitution behaviors.

Figure 2 demonstrates starkly that the rising price change dispersion documented across barcodes in Figure 1 does translate to rising dispersion of household Laspeyres indices, and that households' substitution behaviors do little to draw these distributions closer to mean inflation. It does this using a selection of four product modules, but the conclusions will later be shown to hold more broadly. For each module, a cluster of box and whisker plots summarizes the variation across households in Laspeyres and Paasche measures of category-level inflation.

Solid (hashed) bars reflect the interquartile ranges of household Laspeyres (Paasche) Q4/Q4 indices ending in 2019, 2021, 2022, and 2023. The rising height of solid bars through 2022 captures the sharp rise in item-level price change dispersion from Figure 1 manifesting as a wider range of category-level inflation experiences across households. Comparing the Laspeyres and Paasche bars for a given period reveals how little households' substitution behaviors offset this phenomenon.

The lack of Paasche substitution raises a natural question about whether households truly have access to substitution opportunities. Possibly households in certain areas or shopping at particular retailers have limited access to products exhibiting smaller price changes. To test this possibility, I estimate two measures of households' available inflation rates for each product module. A simple measure checking for location-based variability computes the average inflation rate in that household's market area for the given product module. Another measure computes each module's average inflation rate at the retailer where a household purchases most of its goods in that module. Figure 3 plots household inflation rates against the available inflation measures.

Two pieces of evidence imply that available inflation rates explain little of the households' inflation experiences. Available inflation rates vary much less widely than household inflation rates, implying that most households have access to products exhibiting less extreme price changes

[15] Note that product modules can change rank/position across periods of time. Over the period from 2019-2023, module rankings across quarters exhibit a correlation of 0.58.

[16] Sangani (2023) documents that firms frequently pass on cost shocks in absolute terms, rather than proportionately. That behavior is not responsible for the patterns observed in Figure 1: the interquartile ranges for absolute—rather than log—price changes behave similarly.

in their area and even at their retailer. More directly, regressions of household inflation rates on each measure provide direct checks on the explanatory power of available inflation; each explains only about 10% of variation in household inflation.

C. Household-level Inflation

Figure 4 (also summarized in Table 3) delivers the core result of this paper: the distribution of household inflation rates. The dotted green line, which represents the density of household grocery inflation experiences in the low-inflation period of 2019, illustrates that inflation rates vary widely across households. As the dashed blue and solid red lines indicate, the higher inflation rates of 2021 and 2022 coincided with increasing inflation disparities across households. A dash-dotted yellow line shows that in 2023, when grocery inflation returned to lower levels, the distribution of inflation rates also become less dispersed. The 2022 experience is striking: households in the top decile experienced grocery inflation 7.4 percentage points higher than those in the bottom decile. The results make clear that meaningful differences exist across households based on the specific items they purchase, even when abstracting from differences in prices paid. As the remainder of the paper contends, such differences may be among the *most* consequential determinants of households' varied inflation experiences.

The cross-sectional differences these forces generate are large relative to the time series variation of grocery inflation. Figure 5 plots the time series of PCE Food inflation, which exhibited a similar 1.2 percentage point interquartile range from 2010-2019.[17] Figure 5 also plots the time series properties of the baseline measure as constructed using NielsenIQ data. Its mean reassuringly tracks the PCE Food price index. Cross-sectional variation in grocery inflation, shown as the standard deviation of inflation across households within the quarter, was fairly stable from 2012-2019. A large increase in the standard deviation of inflation in 2022 drives a high correlation of 0.85 between average inflation and its cross-sectional dispersion. Excluding 2022, the mean and standard deviation of grocery inflation exhibit a correlation of 0.36.

Table 3 summarizes the distributions plotted in Figure 4 and provides a deeper history. Its "One price" column displays the baseline results and confirms that, through 2019, interquartile ranges of inflation ranged from 1.2 to 1.6 percentage points according to this measure. Other columns of Table 3 summarize alternative measures that do permit varying degrees of price heterogeneity. Its leftmost column replicates the approach of Kaplan and Schulhofer-Wohl (2017), using only households' price reports and computing inflation among matched goods. The interquartile range of 6.8-7.2 percentage points resembles the range Kaplan and Schulhofer-Wohl recover, which varied from 6.5-8.5 percentage points from 2004 to 2013. By this measure, the interquartile range of inflation experiences in 2022 spanned 12.2 percentage points and a 24.6 percentage point gap separates the 90th and 10th percentiles.[18] Two other measures provide intermediate assumptions.

[17] Food constitutes approximately two-thirds of grocery expenditures.

[18] Two features contribute to a wider distribution of inflation outcomes for this measure: its use of household-level prices instead of national average prices and its consideration of a smaller share of spending, which makes extreme outcomes more likely. Figure A.5 in the Supplemental Appendix decomposes differences between the baseline measure

The "Household + best alternative" column applies household prices when reported, but fills in price changes for non-matched goods using the most relevant available price change. Where possible, these come from the store or retailer where the good was purchased, or else the average price in this household's market area. In the "Best alternative excl. household," price changes are prioritized similarly, but household price reports are skipped. This measure allows prices to differ systematically based on where the household lives and shops, but attempts to insulate the measure from the volatility of individual price reports (driven, for instance, by factors like the timing of sales).[19] On the whole, these results reinforce the message of Kaplan and Schulhofer-Wohl (2017): household-level price heterogeneity is a major additional contributor to inflation heterogeneity. At the same time, the measures indicate that a minority share of household price variation reflects systematic differences across outlets or geography.

That the heterogeneity recovered by tracking detailed item and price information correlates little with demographics challenges an important conclusion the literature drew from Kaplan and Schulhofer-Wohl (2017). Most of the heterogeneity they document arises from the application of household-level price reports. That fact has been cited as evidence that, "an important priority for statistical agencies should be to obtain granular data in as many sectors as possible to keep track of income group-specific expenditure shares, prices paid, and changes in product variety" Jaravel (2021, p. 620-621). Of course, the investment in tracking prices by income group only improves measurement of group-specific inflation to the extent that income groups capture important portions of the variation in prices paid. This paper finds that tracking household prices improves understanding of the extent of within-group inflation heterogeneity; a comparatively small portion of that heterogeneity would pass through to improved group mean differences.

D. Within- and Across-Group Inflation Differences

Table 3 describes a scale of inflation disparities which considerably exceeds those highlighted in studies of inflation differences across groups. This section demonstrates that most studies have captured only a small portion of inflation heterogeneity due to their focus on differences across demographic groups, but not across households within them. Just as aggregating across individual products within a category attenuates the extent of price change variation, aggregating expenditure shares across consumption baskets masks important variation across households.

Figure 6 illustrates the scale of within- and across-group inflation heterogeneity in 2023. Each cell of that figure distinguishes households according to a different demographic grouping—income, age, etc.—and separately plots the distribution of grocery inflation for each group. All groupings convey the same story: so much variation exists within a given demographic group as to make differences across groups comparatively small.

and the Kaplan and Schulhofer-Wohl approach. A focus on matched goods increases measured dispersion even under common prices, but most heterogeneity in the Kaplan and Schulhofer-Wohl measure reflects price differences.

[19] Figures A.1 - A.3 in the Supplemental Appendix display these same distributions, such well as estimates which assume all households face their local average prices. Details for the construction of these estimates can be found in section B of that appendix.

A regression of inflation on household characteristics formalizes the point. Table 4 summarizes three such regressions, each adding new characteristics. The dependent variable in each specification is a household's four-quarter Tornqvist inflation rate using national average price changes, demeaned of time series variation to maintain a cross-sectional focus. The "Characteristics" column employs indicator variables for the listed demographic groups, and the "... with DMA x Quarter Fixed Effects" column adds an interaction between indicator variables for each quarter and each of the 50 best-populated Designated Market Areas in the Nielsen data. Neither specification yields a predictive model, consistent with the impression from Figure 6 that demographic group captures a small portion of the cross-sectional variation in household inflation. Adding the household's lagged inflation rate as a regressor improves predictive power more than any demographic variable.

Table 4 highlights the relationship between this paper's findings and earlier results. Differences across demographic groups are of a similar scale to previous estimates. The monotonically decreasing coefficients for higher income groups are consistent with the findings of Jaravel (2021) and Argente and Lee (2021), albeit smaller than the 0.35 percentage point estimates they recover for an earlier period.[20]

It must be emphasized that this paper does not aim to revisit or diminish the importance of inflation differences by demographic group. Its findings are purely complementary to those careful analyses. As Jaravel notes, the scale of such differences is large—similar in size to the Boskin Commission's estimate of substitution bias—and warrants attention when translating facts about nominal income into statements about well-being.[21] The fact that within-group differences far exceed across-group differences raises separate questions. The remainder of this paper tackles two of these: what generates such large differences in inflation across households, and why do those differences matter?

E. Why So Much Within-Group Inflation Heterogeneity?

Figure 6 illustrates that systematic demographic differences explain a limited portion of the cross-sectional variation in household inflation. This section expands on a core point of the paper: households' grocery bundles vary for many asystematic reasons, and the bundle of a representative agent—even one representing a specific group—ignores a large majority of this variation. I argue that idiosyncratic household demands for specific product varieties, in combination with dispersed price changes across varieties, generates the majority of observed household inflation variation.

To formalize the point, define a price index for the representative agent of group d

$$ln(\bar{\Psi}^d_{t-j,t}) = \sum_{k \in \{\Omega_{t-j} \cup \Omega_t\}} \left(\frac{\bar{s}^d_{k,t-j} + \bar{s}^d_{k,t}}{2} \right) \Delta ln(\bar{p}_{k,t-j,t}) \tag{10}$$

[20] The exclusion of general merchandise, which includes reports of gasoline purchases, likely contributes to a smaller difference across groups.

[21] Boskin et al. (1996, p. 69) believe lower-level substitution bias across individual varieties in a category contributed 0.25 percentage points annually and upper-level substitution across categories contributed 0.15 percentage points annually to growth in the Consumer Price Index.

where $\bar{s}^d_{k,t}$ represents the average expenditure share of product k across households in d.

Imposing that the law of one price holds for all goods simplifies comparisons of inflation rates between groups or within-group differences across households.

$$(11) \qquad ln\left(\bar{\Psi}^d_{h,t-j,t}\right) - ln\left(\bar{\Psi}_{t-j,t}\right) = \sum_{k\in\{\Omega_{t-j}\cup\Omega_t\}} \left(\frac{\bar{s}^d_{k,t-j}+\bar{s}^d_{k,t}}{2} - \frac{\bar{s}_{k,t-j}+\bar{s}_{k,t}}{2}\right)\Delta ln(\bar{p}_{k,t-j,t})$$

$$(12) \qquad ln\left(\Psi_{h,t-j,t}\right) - ln\left(\bar{\Psi}^d_{t-j,t}\right) = \sum_{k\in\{\Omega_{t-j}\cup\Omega_t\}} \left(\frac{s_{h,k,t-j}+s_{h,k,t}}{2} - \frac{\bar{s}^d_{k,t-j}+\bar{s}^d_{k,t}}{2}\right)\Delta ln(\bar{p}_{k,t-j,t})$$

These equations make plain why demographic group comparisons mask the majority of inflation heterogeneity across households. Representative agents' bundles, employing the average expenditure share for each product across many households, are implicitly diversified. Milk may have a minimal expenditure share for a single, lactose-intolerant person but a larger expenditure share for a family with several student athletes. While some demographic groupings (e.g. household size) may capture these differences, they will miss a potentially more important source of variation. The representative agent splits its milk purchases across many varieties, whereas any given household purchases only one variety.

This idiosyncratic variation leads households to have much higher expenditure shares than the representative consumer for the particular varieties they purchase. Equation 12 shows why inflation differences across households require such share differences.

Equation 12 also shows that two other conditions must be met for initial-period expenditure shares to give rise to heterogeneous inflation rates. First, there would be little heterogeneity if price changes $\Delta ln(\bar{p}_{k,t-j,t})$ were similar across goods. Second, even in the presence of heterogeneous price changes, households' substitution away from goods receiving large relative price shocks may mitigate the effects of idiosyncratic exposure.

Figures 1 and 2 illustrated the role of each force. Figure 1 shows that $\Delta ln(\bar{p}_{k,t-j,t})$ can vary widely across close substitutes, and that this potential for households to encounter different price changes grew considerably after 2020. Figure 2 demonstrated the effect of this phenomenon on household inflation rates within narrow product categories, and also revealed that households engage in less substitution than might be expected in response to growing price change heterogeneity within product modules.

Formally, these forces can also be used to decompose households' inflation rates:

$$(13) \qquad ln(\Psi_{h,t-j,t}) = \sum_{k\in\{\Omega_{h,t-j}\cup\Omega_{h,t}\}} \underbrace{s_{h,k,t-j}\Delta ln(\bar{p}_{k,t-j,t})}_{\text{Basket effect}} + \underbrace{\left(\frac{s_{h,k,t}-s_{h,k,t-j}}{2}\right)\Delta ln(\bar{p}_{k,t-j,t})}_{\text{Substitution effect}}.$$

The basket effect in this decomposition represents households' Laspeyres indices. It therefore captures variation arising across households due to the different price changes households face based on the goods in their baskets. The substitution effect is given by the difference between households' Tornqvist and Laspeyres indices. This decomposition can be used to assess the share of total variance in household inflation rates resulting from differing baskets as well as the portion of this variation that households' substitution behaviors offset.

$$(14) \qquad Var(\Psi_{h,t-j,t}) = Var(\text{basket}) + Var(\text{substitution}) + 2\,Cov(\text{basket}, \text{substitution}).$$

Table 5 presents this decomposition for Q4/Q4 inflation rates from 2012-2023. It shows that the basket effect accounts for 97% of cross-sectional variation in household inflation over this period. A negative correlation between households' basket and substitution forces does imply that households facing higher prices tend to be the ones lowering inflation more through substitution. The magnitude of that effect, however, is not large enough to overcome the scale of variation arising from the different price changes affecting household baskets.[22]

V. Model of Household Consumption

Section E points to an important role for households' idiosyncratic commitments to specific product varieties in driving inflation heterogeneity. Evaluating the costs of such inflation requires a model where households exhibit differentiated bundles. McFadden (1974) describes such a model, which Anderson et al. (1992) formally connect to the commonly assumed constant elasticity of substitution (CES) representative consumer. This paper resembles Handbury (2021) in giving households a linear idiosyncratic preference structure within product modules and Cobb-Douglas preferences across modules.[23]

Faced with a collection of varieties yielding different utilities, households find it optimal to purchase one favored product from a given module. Under empirically relevant aggregate demand elasticities, some households have strong preferences for their favored products and are willing to tolerate large changes in relative prices before substituting to a cheaper variety. Faced with the rising spread of price changes documented in Figures 1 and 2, the model generates large and widening inflation disparities for household inflation in each module. Thus adding one feature of realism to an otherwise standard model of consumption endogenizes household inflation experiences. More importantly, the model provides a welfare interpretation for idiosyncratic inflation.

To contextualize the budget shocks this model implies, I embed its grocery consumption block in an intertemporal model of total consumption. Since these features deliver no additional intuition, standard assumptions make the remainder of the model straightforward and set up a simple multi-

[22]Section B.2 in the Supplemental Appendix provides another assessment of the role for substitution in households' overall inflation rates. Figure A.6 plots the Laspeyres and Paasche inflation distributions in place of the Tornqvist indices shown in Figure 4. It reveals that their distributions differ little, reinforcing the conclusion that households with above-average Laspeyres indices do not materially reduce their inflation rates through substitution.

[23]In an early draft, Handbury explicitly estimated the upper-tier elasticity of substitution within a more general CES framework. The results, close to unity, are consistent with the Cobb-Douglas assumption.

stage budgeting problem. Households first make an intertemporal decision about how to allocate lifetime resources based on a forecast of their future inflation rates. Knowing how much they want to spend in each period, households allocate fixed expenditure shares to groceries and to each product module according to nested Cobb-Douglas preferences. Finally, they choose individual products to purchase in each module. This section first describes the allocation of a given budget—the last stages of the household problem—before discussing the grocery and intertemporal allocations.

A. Grocery Budget Allocation (Stages 3-4)

A Cobb-Douglas aggregate across product modules m defines the grocery composite

$$\mathcal{C}_{Gh} = \prod_{m \in G} \left[u_{hm}(c_{hm})\right]^{\lambda_{hm}}. \tag{15}$$

Each household h differs in its desired module expenditure shares λ_{hm}, generating one avenue for heterogeneity across categories. Households with only vegan individuals are unaffected by the price of meat and those with infants are most affected by the price of baby food.

Crucially, the model also features idiosyncratic preferences over individual varieties within product groups. Indexing the varieties in each module by $k \in \Omega_m$, utility from m is

$$u_{hm}(c_{hm}) = \sum_{k \in \Omega_m} c_{hmk} e^{\varphi_{mk} + \varepsilon_{hmk}}. \tag{16}$$

Notationally, real consumption c_{hmk} represents units purchased at prices p_{mk}. The quality of a good φ_{mk} governs the average valuation of product mk across consumers while idiosyncratic preferences are given by ε_{hmk}. The linear indifference curves of this system lead agents to corner solutions. Household h optimally purchases only one good k within each product category m.[24] Formally, they optimize by selecting a quantity of each variety k according to

$$c_{hmk} = \begin{cases} \left(\frac{\lambda_{hm} X_{Gh}}{p_{mk}}\right) \text{ if } k = \underset{k \in m}{\mathbf{argmax}} \left\{\frac{e^{\varphi_{mk} + \varepsilon_{hmk}}}{p_{mk}}\right\} \\ 0 \text{ otherwise.} \end{cases} \tag{17}$$

This mechanism delivers the key intuition of the model: that households do not consume an aggregate bundle of all goods. Instead, each consumes a limited number of individual products. Figures 7 and 8 illustrate this lower- and upper-tier preference structure for two households, one in each panel. Figure 7 depicts households' choices between coffee varieties subject to the preferences given by Equation 17. Blue lines depict the linear indifference curves of that system, the slopes of which $\left(-e^{(\varphi_{ma}+\varepsilon_{hma})}/e^{(\varphi_{mb}+\varepsilon_{hmb})}\right)$ reflect the strength of each household's idiosyncratic preferences over Coffee A and Coffee B. Household 1 strongly prefers Coffee B, while Household 2 has a

[24] Some households purchase more than one type of product in a module during a quarter, which complicates neither theory nor empirics. Such households are viewed to be roughly indifferent between those options. Empirically, the expenditure-weighted average price change across goods determines module-level inflation.

moderate preference for Coffee B. Red dashed and purple dotted lines show relative relative prices in periods $t-1$ and t, including a large increase in the price of Coffee B. That price change does not affect Household 2 but has a large effect on Household 1, who continues to purchase Coffee B after the price increase.

Figure 8 presents the dynamics of this same price change in the upper Cobb-Douglas utility tier across product categories. Blue lines again depict indifference curves, this time between the categories coffee and milk. Again the dashed red and dotted purple lines depict relative prices in period $t-1$ and t, respectively. While Household 2 experiences no change in the relative price of these two categories, Household 1 faces a significantly higher cost of coffee. It responds by reducing the quantity of coffee purchased and moves to a less gratifying utility level depicted by the gray indifference curve.

The two households depicted in Figures 7 and 8 experience widely differing inflation outcomes due to their exposure to item-level variation in price changes. A representative agent approach applying the same Cobb-Douglas preferences depicted in Figure 8 would expose the representative household to the comparatively stable *average* relative prices of coffee and milk. Approaches constructing demographic group consumption baskets or applying CPI category prices to individual households' bundles would similarly fail to capture the within-category price change variation across items generating this heterogeneity.

Assuming agents draw their idiosyncratic preferences for goods from a Type I extreme value distribution $\Gamma()$ with a scale parameter of 0 and category-specific shape parameter of $1/(\sigma_m - 1)$ implies helpful aggregation properties. Anderson et al. (1992, p. 85-90) demonstrate that with $\varepsilon_{hmk} \sim \Gamma(0, \frac{1}{(\sigma_m - 1)})$, demands from this system correspond to those of a CES representative consumer having an elasticity of substitution σ_m and module utility:

$$\tilde{u}_m = \left[\sum_{k \in \Omega_m} \left(c_{mk} e^{\gamma_{mk}} \right)^{\frac{\sigma_m - 1}{\sigma_m}} \right]^{\frac{\sigma_m}{\sigma_m - 1}}. \tag{18}$$

Argente and Lee estimate elasticities for each NielsenIQ product module (2021, p. 925). The median elasticity across product modules has $\sigma_m \approx 20$, and half of the product modules feature an elasticity between 10 and 40.[25] Table 6 summarizes a distribution of simulated ε_{hmk} under this range of empirically relevant elasticities. It demonstrates that these assumptions result in substantial item-level commitments for some households. For a product category with an elasticity of 20, ten percent of households value their favorite good at least 34% more than the typical good of similar quality.[26] Substitution rates differ only slightly across income groups, with low-income

[25]Note that these estimates are within homogeneous NielsenIQ product modules (e.g. frozen chicken). Other benchmark estimates tend to be estimated over product groups (e.g. frozen meat, poultry, and seafood) and range from 3-7 (see Hottman et al. (2016) or Redding and Weinstein (2020)).

[26]Table A.2 in the Supplemental Appendix reports relative values of the first and second most-preferred goods. At an elasticity of 20, ten percent of households value a unit of their favorite good at least 12% more than their next favorite good. For the quarter of product modules with an elasticity below 10, that figure is 23%.

households exhibiting elasticities approximately 8% higher than high-income households.[27]

Introducing a more realistic setting where households buy only one variety from a given product category generates a different inflation landscape, one subject to the full idiosyncratic volatility of item-level price changes displayed in Figures 1 and 2. The framework is capable of generating the observed inflation disparities because households endure significant relative price changes without substituting to an alternative with lower price growth. The next section expands on the intuition just discussed to connect grocery inflation and welfare in this system.

A.1 Deriving Bounds to the Household's Cost Function

For a household with preferences given by this model, the unit price of grocery composite $\mathcal{C}_G$ depends only on the set of goods optimally purchased in each module. It takes the form

$$P_G(\overrightarrow{p}, \overrightarrow{\varepsilon}_h) = \prod_{m \in G} \left[\lambda_{hm}^{-\lambda_{hm}} \min_{k \in \Omega_m} \left(\frac{p_{mk}}{e^{\varphi_{mk} + \varepsilon_{hmk}}} \right) \right]^{\lambda_{hm}}. \tag{19}$$

With this cost function, and denoting good $m\ell$ as the cost-minimizing variety in product category m under prices in t and mk as the cost-minimizing variety under $t-j$ prices, the price index measuring the changing cost of $\mathcal{C}_G$ under prices p_t and p_{t-j} would be

$$\Psi_{h,t-j,t} = \frac{P_G(p_t, \varepsilon_h)}{P_G(p_{t-j}, \varepsilon_h)} = \prod_{m \in G} \left[\frac{\frac{p_{m\ell t}}{e^{\varphi_{m\ell} + \varepsilon_{hm\ell}}}}{\frac{p_{mkt-j}}{e^{\varphi_{mk} + \varepsilon_{hmk}}}} \right]^{\lambda_{hm}} = \prod_{m \in G} \left[\left(\frac{p_{m\ell t}}{p_{mkt-j}} \right) \left(\frac{e^{\varphi_{mk} + \varepsilon_{hmk}}}{e^{\varphi_{m\ell} + \varepsilon_{hm\ell}}} \right) \right]^{\lambda_{hm}}. \tag{20}$$

The presence of unobservable idiosyncratic preference terms for every good the household purchases makes direct estimation of this index appear hopeless. I demonstrate below that this cost function can be bounded and approximated by statistical price indices. To see this, first group the index into modules for which $m\ell = mk$ and those for which it does not. The former contains the set of matched goods, denoted $\Omega^*_{h,t-j,t}$. The latter contains all others in the set of goods purchased either period

$$\Psi_{h,t-j,t} = \left[\prod_{m \in \Omega^*_{h,t-j,t}} \left[\left(\frac{p_{mkt}}{p_{mkt-j}} \right) \right]^{\lambda_{hm}} \right] \left[\prod_{n \in \left\{ \Omega_{h,t-j,t} - \Omega^*_{h,t-j,t} \right\}} \left[\left(\frac{p_{n\ell t}}{p_{nkt-j}} \right) \left(\frac{e^{\varphi_{nk} + \varepsilon_{hnk}}}{e^{\varphi_{n\ell} + \varepsilon_{hn\ell}}} \right) \right]^{\lambda_{hn}} \right]. \tag{21}$$

Among matched goods, appeal terms cancel and the appropriate price index takes an unambiguous and estimable form.[28] The term for unmatched goods is more ambiguous, but a straightforward appeal to theory implies that we can safely replace the unobservable terms. If households optimally

[27] This being the case, the model does not incorporate Handbury's interaction between households' idiosyncratic preference strength and their total spending.

[28] Note that the Cobb-Douglas demand system implies that expenditure shares for each module remain constant. A Laspeyres, Paasche, or Tornqvist index would return identical estimates of this term in such a setting.

choose product nk under period $t-j$ prices and $n\ell$ under period t prices, such that

$$(22) \qquad nk = \underset{i\in\Omega_n}{\operatorname{argmin}} \left(\frac{p_{nit-j}}{e^{\varphi_{ni}+\varepsilon_{hni}}}\right) \text{ and } n\ell = \underset{i\in\Omega_n}{\operatorname{argmin}} \left(\frac{p_{nit}}{e^{\varphi_{ni}+\varepsilon_{hni}}}\right)$$

the following may be inferred

$$(23) \qquad \frac{p_{nkt-j}}{e^{\varphi_{nk}+\varepsilon_{hnk}}} \leq \frac{p_{n\ell t-j}}{e^{\varphi_{n\ell}+\varepsilon_{hn\ell}}} \text{ and } \frac{p_{n\ell t}}{e^{\varphi_{n\ell}+\varepsilon_{hn\ell}}} \leq \frac{p_{nkt}}{e^{\varphi_{nk}+\varepsilon_{hnk}}}.$$

From these relationships it can readily be seen

$$(24) \qquad \frac{p_{n\ell t}}{p_{n\ell t-j}} \leq \left(\frac{p_{n\ell t}}{p_{nkt-j}}\right)\left(\frac{e^{\varphi_{nk}+\varepsilon_{hnk}}}{e^{\varphi_{n\ell}+\varepsilon_{hn\ell}}}\right) \text{ and } \left(\frac{p_{n\ell t}}{p_{nkt-j}}\right)\left(\frac{e^{\varphi_{nk}+\varepsilon_{hnk}}}{e^{\varphi_{n\ell}+\varepsilon_{hn\ell}}}\right) \leq \frac{p_{nkt}}{p_{nkt-j}}.$$

Thus for each product the household substituted between, optimal decision-making implies observable bounds on its contribution to the household's exact price index. Using the price change of good $n\ell$ (corresponding to a Paasche index) bounds that contribution from below, while the price change of good nk (corresponding to a Laspeyres index) bounds the contribution from above. Headline results reported throughout the paper, by utilizing a Tornqvist index, provide an estimate which averages between these two bounds. Figure A.6 reveals the inflation distributions to be similar across all three measures.[29]

A.2 Household Inflation as a Budget Shock

The welfare implications of idiosyncratic inflation become clearer when translated into annualized budget shocks. Table A.1 in the Supplemental Appendix summarizes spending in the 2019 Consumer Expenditure Survey, when Americans spent an average of \$6,688 on groceries. The median inflation experience from 2019Q4-2020Q4 implies annualized grocery spending of \$6,902 needed to achieve the same level of satisfaction. Such an inflation experience can be thought of as a \$214 change in the annual cost of the household's groceries over that period.

Table 7 summarizes the distribution of such shocks. It shows that 10% of households instead received grocery budget shocks of \$334 or more, while another 10% received grocery budget shocks of \$107 or less. This \$227 gap between the grocery budget shocks households experience at the 90th and 10th percentiles of inflation grew to \$572 in 2022, highlighting the cost of a wider inflation distribution. Even if average wages and returns to savings were perfectly indexed to compensate for average aggregate inflation, some households face significantly higher cost shocks than this indexation offsets.

B. Grocery Share of Total Consumption (Stage 2)

The dollar-denominated shocks of Table 7 show how extreme grocery inflation can affect some households. Yet these shocks may be better understood by contextualizing them within households'

[29] Supplemental Appendix B.2 discusses the atypical relation between Paasche and Laspeyres after 2021.

total expenditures and accounting for how changing expenditures affect well-being. As a simple means of adding that context, assume that households allocate some share of spending γ_h to groceries and the remainder to an outside good $\mathcal{C}_N$.[30,31]

$$(25) \qquad \mathcal{C}_{ht} = (\mathcal{C}_{Ght})^{\gamma_h} (\mathcal{C}_{Nht})^{1-\gamma_h}$$

Under this system, the household's total expenditures $X_{ht} = X_{Ght} + X_{Nht}$ relate to a unit of consumption $\mathcal{C}_h$ through the grocery and non-grocery cost functions P_{Ght} (as in Equation 19) and P_{Nht}, according to

$$(26) \qquad X_{ht} = P_{ht}\mathcal{C}_{ht} \quad \text{where}$$

$$(27) \qquad P_{ht} = \left(\frac{P_{Ght}}{\gamma_h}\right)^{\gamma_h} \left(\frac{P_{Nht}}{1-\gamma_h}\right)^{1-\gamma_h}.$$

As a straightforward benchmark for scaling the implications of heterogeneous grocery inflation, consider a scenario with zero non-grocery inflation: $P_{Nht-j} = P_{Nht}$. The 11.4 percent median grocery inflation experience from 2021Q4 to 2022Q4 translates to a 1.4 percent increase in the household's overall cost of living. Households experiencing 90th percentile inflation of 15.2 percent instead see a 1.8 percent increase in the overall cost of living while those experiencing 10th percentile inflation see only a 0.9 percent increase. Absent any inflation heterogeneity in other spending categories, such disparate grocery inflation experiences generate cost of living gaps of 0.9 percentage points.

Taking the log change in Equation 26 defines an important identity

$$(28) \qquad \Delta ln(P_{ht}) = \Delta ln(X_{ht}) - \Delta ln(\mathcal{C}_{ht}).$$

Households have two available responses to inflation: increasing expenditures or reducing consumption. Without increasing expenditures, a household facing a 1.8 percent increase in its cost of living must reduce real consumption by an equal amount. Assuming a specific utility function $\nu()$ over total consumption provides an estimate of how total well-being responds to such a decline.

The class of Constant Relative Risk Aversion (CRRA) functions provides natural candidates for $\nu()$. Under these functional forms, the coefficient of relative risk tolerance θ governs the relationship

[30] The base-period price of this outside good, $p_{N,0} \equiv 1$, may be taken as the system's numeraire.

[31] In the Consumer Expenditure Survey, households' grocery purchases range from 12.4-13.7% of total consumption—which excludes the "Personal insurance and pensions" and "Cash contributions" categories in Table A.1—over the 2018-2023 period.

between consumption and utility.[32] CRRA utility takes the form

$$\nu(\mathcal{C}_{ht}) = \frac{\mathcal{C}_{ht}^{1-\frac{1}{\theta}} - 1}{1 - \frac{1}{\theta}}. \tag{29}$$

With $\theta = 0.25$, as estimated by Barsky et al. (1997), a 15% reduction in real consumption corresponds to a 20% loss of utility whereas a 15% increase in real consumption corresponds to about a 12% increase in utility. This concavity governs the magnitude of households' welfare losses under extreme inflation outcomes.[33]

Households whose cost of living rises faster than their wage might not offset the entirety of their budget shock by reducing consumption. They may be able to increase expenditures at the cost of reducing future consumption. To capture these dynamics requires adding an intertemporal dimension to the model.

C. Allocating Expenditures Across Time (Stage 1)

Having seen how households would allocate a given budget under the assumed preference structure, it remains to establish how households decide their expenditures each period. Begin by assuming that households value the stream of utilities $\nu(C_{h\tau})$ from periods $\tau = t, t+1, ..., T$ as a discounted sum where β is the discount rate

$$\mathcal{V}\left(\{\mathcal{C}_{h\tau}\}_{\tau=t}^{T}\right) = \sum_{\tau=t}^{T} \beta^{\tau-t} \nu(\mathcal{C}_{h\tau}). \tag{30}$$

In this framework, θ governs both the marginal value of consumption within a given period and the elasticity of substitution across periods.[34] A value of $\theta = 0.25$ implies that a household preferring equal consumption between adjacent periods would need to consume 16% more in the next period to compensate for a 10% consumption reduction in the current period.

The full objective of each household is to plan a path of consumption $\{\mathcal{C}_{h\tau}\}_{\tau=t}^{T}$ and savings $\{A_{h\tau}\}_{\tau=t}^{T}$ to maximize expected lifetime utility subject to a flow of budget constraints:

$$\max_{\{\mathcal{C}_{h\tau}\}_{\tau=t}^{T}} \left\{\mathbb{E}\left[\mathcal{V}\left(\{\mathcal{C}_{h\tau}\}_{\tau=t}^{T}\right)\right]\right\} \equiv \max_{\{X_{h\tau}\}_{t=\tau}^{T}} \left\{\mathbb{E}\left[\mathcal{V}\left(\left\{\frac{X_{h\tau}}{P_{h\tau}}\right\}_{\tau=t}^{T}\right)\right]\right\}$$

$$\text{s.t.} \quad A_{h\tau+1} = A_{h\tau}(1 + r_\tau) + Y_{h\tau} - X_{h\tau} \;\; \forall\, \tau \quad \text{and} \quad A_{hT} \geq 0. \tag{31}$$

Households attempting to plan a consumption path need to form expectations over many uncer-

[32] These same preferences are commonly characterized through the coefficient of relative risk aversion $\frac{1}{\theta}$.

[33] In the Supplemental Appendix, the left panel of Figure A.10 illustrates such preferences under differing θ values. Figure A.11 uses Equation 29 to translate inflation to changes in flow utility assuming households do not change their grocery budgets and face no inflation in non-grocery categories.

[34] The right panel of Figure A.10 in the Supplemental Appendix plots indifference curves between consumption in two adjacent quarters under a range of elasticities of intertemporal substitution θ.

tain quantities: grocery and non-grocery inflation, income changes, and uncertain asset returns. To maintain focus on the role of idiosyncratic grocery inflation, assume households face no uncertainty over income, real interest rates, or aggregate inflation. Specifically, households know their income path $\{Y_{h\tau}\}_{\tau=t}^{T}$ and that path automatically adjusts to compensate for aggregate inflation $\bar{\Psi}_{\tau}$. For simplicity, assume that aggregate inflation is common to both grocery and non-grocery products: $\Delta ln(\bar{\Psi}_{\tau}) = \Delta ln(P_{G\tau}) = \Delta ln(P_{N\tau})$. Bonds pay interest which varies from a constant real rate r only to offset losses from aggregate inflation: $(1+r_{\tau}) = \bar{\Psi}_{\tau}(1+r)$. Against this backdrop of neutral aggregate inflation, households only face uncertainty over the idiosyncratic costs of their optimal grocery bundle $P_{Gh}(\overrightarrow{p}_{\tau}, \overrightarrow{\varepsilon}_{h\tau})$.

With a typical household encountering thousands of products, even this simplified problem lacks tractability if households forecast individual prices $\overrightarrow{p}_{\tau}$. The problem is simplified, however, if households respond to that challenge by forecasting not individual item-level prices, but future idiosyncratic inflation rates $\mathbb{E}[\{\Psi_{h\tau}\}_{\tau=t+1}^{T}]$. As Deaton and Muellbauer (1980, p. 309-310) discuss, intertemporal consumption problems become amenable to a multi-stage budgeting approach when households' price expectations for individual items move in parallel. Households with such expectations first make a plan for the intertemporal allocation of expenditures based on average expected price levels, then decide which products to purchase once true relative prices are revealed.

Noting that $\mathcal{C}_{h\tau}P_{h\tau} = (1+r_{\tau})A_{h\tau} + Y_{h\tau} - A_{h\tau+1}$, the intertemporal problem can be conceptualized as choosing a plan for current-period consumption $\mathcal{C}_{h\tau}$ and a path of future assets $\{A_{h\tau}\}_{t+1}^{T}$. The household arriving at period t with assets A_{ht} and realized prices $\overrightarrow{p}_{t}$ forms inflation expectations $\mathbb{E}[\{\Psi_{h\tau}\}_{\tau=t+1}^{T}]$ and makes a savings plan satisfying the Euler equation:

$$\nu'\left(\mathcal{C}_{h\tau}\right) = \beta\mathbb{E}\left[\left(\frac{1+r_{t+1}}{\Psi_{ht+1}}\right)\nu'\left(\mathcal{C}_{h,t+1}\right)\right] \tag{32}$$

Idiosyncratic inflation in this framework has similar effects as aggregate inflation in a typical intertemporal problem. Higher price indices reduce the real return on nominal assets arriving in period $t+1$ (first term on the right) and reduce the real consumption value of a given unit of nominal expenditures, raising the marginal utility of future consumption (second term on the right).

In a world with no other uncertainties, the optimal consumption path for a household depends on its expectations about future inflation. Consider now two extreme potential outcomes. In one, the household's above-average inflation rate reflects a temporary drift of prices in its preferred basket from average prices. The household can expect a below-average inflation rate next period so that its cost of living returns to the average level. The Euler equation offers intuitive guidance on how that expectation informs the household's consumption decision. Both because nominal expenditures in $t+1$ yield higher utility and because saved assets provide greater satisfaction in future quarters, the household would spend less in t under this paradigm. In welfare terms, however, the household should not be much affected so long as they have access to credit that allows them to smooth consumption through the period of temporarily high prices.

In the opposite extreme, suppose that the above-average price changes for goods in this household's preferred bundle reflect permanent cost shocks unique to those products. Rather than mean-reverting to an average price level via below-average inflation rates, the household tends to continue exhibiting average inflation. In such a case, the household's above-average inflation rate persists in a higher cost of living $P()$ and a reduction in real assets. In essence, it manifests as a permanent reduction in household wealth. With no expectation of future consumption being comparatively cheaper, this household has no incentive to substitute intertemporally.

VI. Costs of Inflation Heterogeneity

The model laid out in the preceding section provides the tools needed to evaluate the lifetime utility consequences of grocery inflation shocks. For a given utility parameterization and an assumed path for how a household's cost of living evolves following a high idiosyncratic inflation draw, Equation 32 characterizes the household's optimal consumption path and Equation 30 estimates the utility value of that path. Table 8 summarizes the lifetime reduction from a 4% grocery inflation draw under various utility parameterizations and under a simple inflation process of varying persistence. Rows summarize the elasticity of intertemporal substitution θ and columns summarize how long it takes the household to return to an average cost of living under a constant inflation gap. Results vary less across choice of θ than the persistence of the price shocks. Moderately persistent price shocks that leave households' costs of living elevated for 2-3 years imply lifetime utility losses of approximately 0.1% across the range of θ values.

An autoregressive/moving average model of idiosyncratic inflation can help shed light on the crucial question of persistence. Suppose households' one-quarter inflation rates follow an autoregressive or moving average process, as in

$$\tilde{\Psi}_{h,t-1,t} \equiv \Psi_{h,t-1,t} - \bar{\Psi}_{t-1,t} = \sum_{\ell=1}^{L_{AR}} \phi_{-\ell} \tilde{\Psi}_{h,t-\ell-1,t-\ell} + \sum_{\ell=1}^{L_{MA}} \delta_{-\ell} \epsilon_{h,t-\ell-1,t-\ell} + \epsilon_{h,t-1,t}. \tag{33}$$

If households' idiosyncratic cost of living featured perfect mean reversion, we would expect those with above average inflation rates last period (high $\tilde{\Psi}_{h,t-2,t-1}$) to exhibit *below-average* inflation rates this period (low $\tilde{\Psi}_{h,t-1,t}$). A first-order moving average model of that process would recover an estimate of $\delta_{-1} \approx -1$. To illustrate this in a concrete two-period example, suppose aggregate inflation of 1% in two successive periods raises the aggregate cost of living index from 100 to 101 in the first period, then to 102 in the second. A household experiencing 3% inflation in the first period (2 percentage points above average) has a cost of living index of 103. To return to the average cost of living in the subsequent period, this household needs an inflation rate of -0.97% (approximately 2 percentage points below average). If instead it experiences average inflation of 1%, it will continue to suffer an above-average cost of living (in this case 104).

In the scenario with a strong negative serial correlation, households drift briefly from common cost of living trends. In the opposite extreme, a process which exhibits no serial correlation between inflation rates would imply that households' idiosyncratic costs of living follow a random walk. In any intermediate case, high relative inflation might be followed by a sequence of below-average inflation rates over a horizon of L periods to bring households closer to the average cost of living.

Table 9 summarizes the estimates for several variants of Equation 33. Results confidently reject the idea that high relative inflation quickly offsets in the subsequent periods. The degree of persistence implied by the point estimates is more consistent with idiosyncratic inflation moving households to persistently higher cost of living paths. The first-order moving average model implies that a household experiencing inflation 2 percentage points above average in one period should expect an inflation rate just 0.18 percentage points below average in the subsequent period.[35,36]

To evaluate whether household inflation experiences mean revert at longer horizons, consider household-level price indices which cumulate individual inflation draws. Whether and how quickly the cross-sectional variance of these cumulative inflation rates grows informs the degree of mean reversion exhibited over moderate horizons. If household inflation exhibited no mean reversion (i.e. followed a random walk) the cross-sectional variance of cumulative inflation would grow as a linear multiple of the one-quarter variance of inflation rates. If household inflation rates are mean reverting to any degree, the cross-sectional variance of cumulative inflation grows more slowly.

Campbell et al. (2012, p. 48-57) describe an intuitive metric leveraging this logic to assess the relative permanence of price changes: the variance ratio. Formally, define the variance ratio at a horizon of k quarters as

$$VR(k) = \frac{Var(\Psi_{h,t-k,t})}{kVar(\Psi_{h,t-1,t})} \tag{34}$$

where $\Psi_{h,t-k,t}$ chains households' one-quarter Tornqvist inflation rates between periods $t-k$ and t. Figure 9 plots the estimated variance ratio up to a 12-quarter horizon. Consistent with the slightly negative MA(1) coefficient estimated for equation 33, the variance ratio of 0.89 at a two-quarter horizon is less than unity but much larger than the value of $VR(2) = 0.5$ that would result from perfect mean reversion. The variance ratio continues to decline over four quarters to a value of 0.74 and converges to a value around 0.71.

The value to which this variance ratio converges is informative about the persistence of relative price changes for households' baskets. Imagine that a household's inflation rate is determined in part by a stationary process whose influence ultimately mean reverts $\varphi_{h,\tau-1,\tau}$ and in part by a random component $\psi_{h,t-k,t}$

$$\Psi_{h,t-k,t} = \prod_{\tau=t-k}^{t} \Psi_{h,\tau-1,\tau} = \prod_{\tau=t-k}^{t} (\varphi_{\tau-1,\tau})(\psi_{h,\tau-1,\tau}). \tag{35}$$

Under such a partition, Campbell et al. (2012, p. 57) note that the variance ratio's long-run value captures the relative variation of the transitory and permanent components of household inflation

[35] A moving average model may be more naturally written as $\Psi_{h,t-1,t} = \alpha_0\bar{\Psi}_{t-1,t} + \delta_{-1}\epsilon_{h,t-1,t}$ without constraining $\alpha_0 = 1$ as in Equation 33. Estimated in this way, $\alpha_0 = 0.98$ leaving δ_{-1} unchanged.

[36] The results summarized in Table 9 are little changed by excluding the post-pandemic period.

rates. In particular:

$$VR(k) \to 1 - \frac{Var(\varphi_{h,t-1,t})}{Var(\psi_{h,t-1,t})} \tag{36}$$

A long-run variance ratio of 0.71 is consistent with the random inflation component exhibiting 2.5 times greater volatility than the mean-reverting component. Consequently, large portions of the idiosyncratic price changes households encounter should be expected to accumulate rather than serially offset.

Figure 10 uses the intertemporal model of Equation 31 to provide an estimate of the reductions in lifetime utility that result from each draw in the empirical distribution of grocery inflation. It does this under the assumption that a household's cost of living reverts to the average cost of living over a period of three years following an idiosyncratic inflation shock.[37] The 2022 distribution of inflation rates imply that households with 90th and 10th percentile inflation outcomes feature lifetime utility reductions that differ by 0.2 percentage points. It is remarkable that such reductions in total utility across *all* product categories result only from grocery inflation.

A. *Household Responses to Idiosyncratic Inflation*

Table 10 summarizes regressions which assess how a household's consumption and shopping behaviors interact with its inflation rate. Values in the table reflect the coefficient of households' four-quarter Tornqvist inflation rates—interacted with an income group dummy given by the table row—in a regression on the variable given by the table column. Each regression includes (not displayed) the same demographic controls appearing in the last two columns in Table 4.

The first two columns assess how households adjust nominal grocery expenditures or implied real consumption.[38] Results imply that households do not pass the entirety of inflation shocks onto spending increases. That behavior might be expected if households' individual inflation rates went unnoticed and they purchased their usual basket of goods regardless of costs. Instead, households appear to offset just under half of the budget shocks inflation presents by reducing consumption. Such a response is consistent with a key prediction of the model: that households perceiving a persistent shock to the price of their grocery bundle would reduce consumption as if having received a wealth shock.

All income groups exhibit this consumption-reducing behavior to slightly varied degrees. Households earning more than \$100k per year appear most willing to reduce grocery consumption in response to grocery inflation, while those in the \$50-100k income group offset only about one third of their implied budget shock with reduced consumption. Shopping behavior margins show com-

[37] The computations assume that a household's idiosyncratic inflation rate runs below average at the constant rate needed to return to the average cost of living within the specified window. This linear path back to an average cost of living features more aggressive mean reversion than implied by the models presented in Table 9. Taken literally, estimates of equation 33 and the variance ratio analysis presented in Figure 9 imply that households *never* fully return to mean inflation. As Table 8 illustrates, such a process would generate larger costs to extreme inflation than those shown in Figure 10.

[38] As inflation links real consumption to nominal spending by identity, (column 1) - (column 2) = 1.

paratively little relation. Households exhibit no discernible response to the number of stores visited or the dollar value of coupons used. The model detects a small response for number of trips households in the middle income group make, shopping slightly more often in periods of rising prices. For households in the $25-50k income bin, the average shopping response to a 90th percentile idiosyncratic inflation shock in 2022 corresponds to one additional shopping trip per quarter.

VII. Conclusion

It should not surprise that beneath the average inflation rate—even for a group of demographically similar households—lies a distribution of experiences. Yet this paper captures important new facts about the scale and interpretation of that distribution, most importantly how it evolved during the recent period of high aggregate inflation.

Differences in household consumption choices over product varieties, in combination with those goods' differing price changes, generate a form of inflation heterogeneity not captured when aggregating across goods or consumers. This heterogeneity is responsible for more variation across households than demographic or behavioral differences that tend to drive narratives of inflation disparities. While it is tempting to assume that the welfare consequences of switching between coffee or milk varieties must be small, this source of heterogeneity should not be dismissed. Some households absorb large price shocks without substituting toward goods with lower price growth, implying an inability or unwillingness to mitigate idiosyncratic inflation. A model endogenizing this behavior using idiosyncratic preferences demonstrates that such inflation experiences have the interpretation of shocks to the household's cost of achieving a given level of well-being.

The heterogeneity I document and framework for evaluating it point to an important new channel for understanding households' perceived costs of inflation. They imply that even a neutral aggregate inflation may have distributional costs, with some households experiencing significant welfare losses. Over the four quarters ending 2022Q4, when price change dispersion almost doubled in many product categories, households in the 90th inflation percentile received annualized budget shocks approximately $300 larger than the median household. The higher cost of living that follows a bad inflation draw tends to persist, rather than mean revert, implying non-trivial lifetime welfare reductions. This mechanism sheds light on why some households report much more economic pain from inflation than the aggregate statistics imply. Those households hate grocery inflation for a simple reason: because it genuinely causes them economic distress.

References

Alvarez, F., M. Beraja, M. Gonzalez-Rozada, and P. A. Neumeyer (2019). From hyperinflation to stable prices: Argentina's evidence on menu cost models. *The Quarterly Journal of Economics 134*(1), 451–505.

Anderson, S. P., A. De Palma, and J.-F. Thisse (1992). *Discrete Choice Theory of Product Differentiation.* MIT press.

Argente, D. and M. Lee (2021). Cost of living inequality during the great recession. *Journal of the European Economic Association 19*(2), 913–952.

Barsky, R. B., F. T. Juster, M. S. Kimball, and M. D. Shapiro (1997). Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study. *The Quarterly Journal of Economics 112*(2), 537–579.

Boskin, M., E. Dulberger, R. Gordon, Z. Griliches, and D. Jorgenson (1996, December). *Toward a More Accurate Measure of the Cost of Living: Final Report to the Senate Finance Committee.* Washington, DC: U. S. Government Printing Office for the U. S. Senate Committee on Finance.

Campbell, J. Y., A. W. Lo, and A. C. MacKinlay (2012). *The Econometrics of Financial Markets.* Princeton, NJ: Princeton University Press.

Cavallo, A. and O. Kryvtsov (2024). Price discounts and cheapflation during the post-pandemic inflation surge. *Journal of Monetary Economics 148*, 103644.

Chen, T., P. Levell, and M. O'Connell (2024). Cheapflation and the rise of inflation inequality. *Institute for Fiscal Studies Working Paper*.

Deaton, A. and J. Muellbauer (1980). *Economics and Consumer Behavior.* Cambridge University Press.

Diewert, W. E. (1976). Exact and superlative index numbers. *Journal of econometrics 4*(2), 115–145.

Georgarakos, D., K. H. Kim, O. Coibion, M. Shim, M. A. Lee, Y. Gorodnichenko, G. Kenny, S. Han, and M. Weber (2025). How costly are business cycle volatility and inflation? a vox populi approach. Technical report, National Bureau of Economic Research.

Handbury, J. (2021). Are poor cities cheap for everyone? non-homotheticity and the cost of living across us cities. *Econometrica 89*(6), 2679–2715.

Hill, R. J. and D. Melser (2008). Hedonic imputation and the price index problem: an application to housing. *Economic Inquiry 46*(4), 593–609.

Hobijn, B. and D. Lagakos (2005). Inflation inequality in the united states. *Review of Income and Wealth 51*(4), 581–606.

Hottman, C. J., S. J. Redding, and D. E. Weinstein (2016). Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics 131*(3), 1291–1364.

Jaravel, X. (2021). Inflation inequality: Measurement, causes, and policy implications. *Annual Review of Economics 13*, 599–629.

Jaravel, X. (2024). Distributional consumer price indices, working paper.

Kaplan, G. and S. Schulhofer-Wohl (2017). Inflation at the household level. *Journal of Monetary Economics 91*, 19–38.

Kostyshyna, O. and M. Ouellet (2024). Household food inflation in canada. Technical report, Bank of Canada Staff Working Paper.

McFadden, D. (1974). *The Measurement of Urban Travel Demand*, Volume 3. North-Holland.

McGranahan, L. and A. L. Paulson (2005). Constructing the chicago fed income based economic index-consumer price index: Inflation experiences by demographic group: 1983-2005, working paper 2005-20.

Nakamura, E., J. Steinsson, P. Sun, and D. Villar (2018). The elusive costs of inflation: Price dispersion during the us great inflation. *The Quarterly Journal of Economics 133*(4), 1933–1980.

Orchard, J. (2020). Household inflation and aggregate inflation, working paper.

Pathak-Chalise, P. (2025). Household-level food price inflation heterogeneity: Evidence and insights from the us consumer panel data (2013-2023), working paper.

Pollak, R. A. (1989). *The Theory of the Cost-of-Living Index.* Oxford University Press, USA.

Redding, S. J. and D. E. Weinstein (2020). Measuring aggregate price indices with taste shocks: Theory and evidence for ces preferences. *The Quarterly Journal of Economics 135*(1), 503–560.

Sangani, K. (2023). Pass-through in levels and the incidence of commodity shock, working paper.

U.S. Bureau of Economic Analysis (2024). Personal income and outlays. Retrieved from FRED, Federal Reserve Bank of St. Louis, January 15, 2025.

U.S. Bureau of Labor Statistics (2019). Consumer expenditure survey. Table R-1.

U.S. Census Bureau (2020). Zip code tabulation areas gazetteer files.

Figure 1
Price Change Dispersion Across Items Within Product Modules

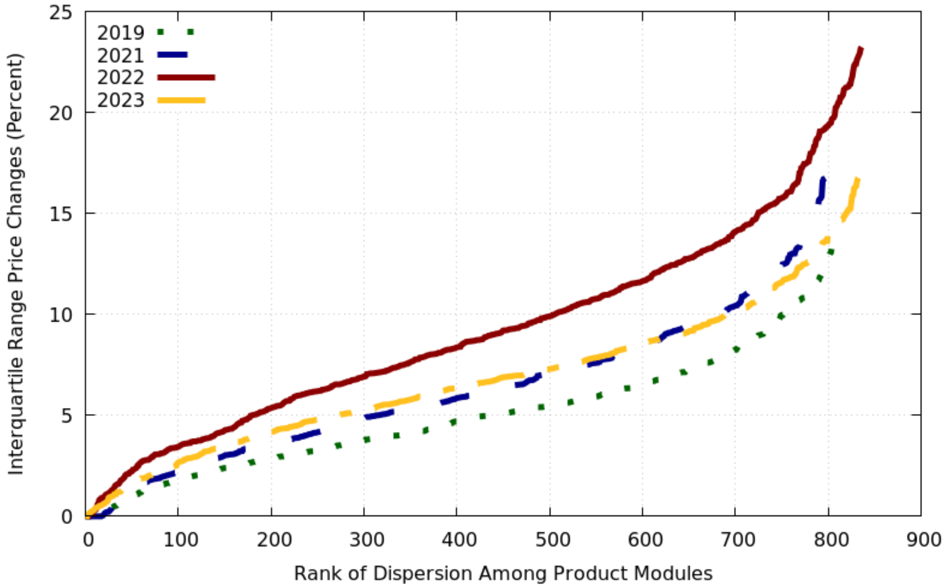

The height of each line reflects the gap between the 75th and 25th percentile of national average log price changes across goods in that module. The figure excludes modules in the general merchandise department and uses only prices recorded in the Retail Scanner data. Each barcode is importance weighted by its share of total expenditures so that each category's interquartile range captures 50% of spending rather than 50% of barcodes. Product modules with interquartile ranges below the 10th or above the 90th percentile excluded to improve visibility.

Figure 2
Spread of Module Inflation Rates

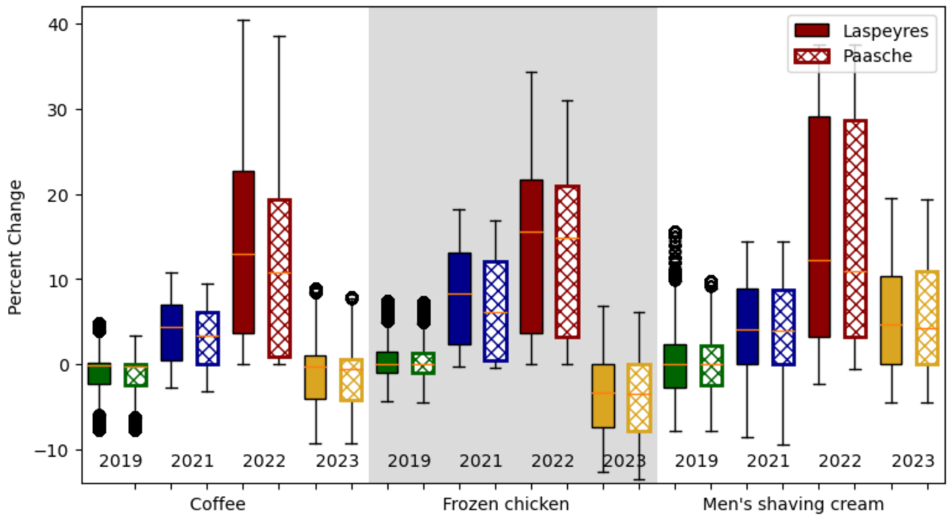

Boxes summarize the interquartile range of household four-quarter inflation rates in the given product modules. Whiskers extend to the remaining variation, with dots representing observations beyond 1.5x the interquartile range. Observations outside of 5th-95th percentiles are excluded to improve visibility. For households buying only one variety within the module in the base period (Q-4), inflation is computed using that good's price change between Q-4 and Q. For a household buying multiple varieties, this is the Q-4 expenditure-weighted average price change across these goods.

Figure 3
Relationship of Household and Available Module Inflation Rates

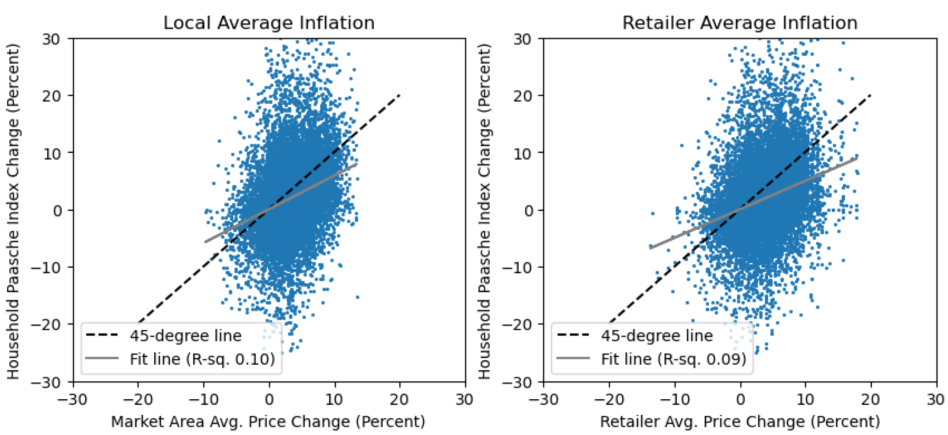

Each dot represents a household's four-quarter average price change within a product module. Position on y-axis reflects household's Paasche-weighted average price change, while position on x-axis shows price changes in same module in household's Nielsen DMA or at the retailer where the household purchases majority of this product.

Figure 4
Distribution of Household Grocery Inflation Experiences

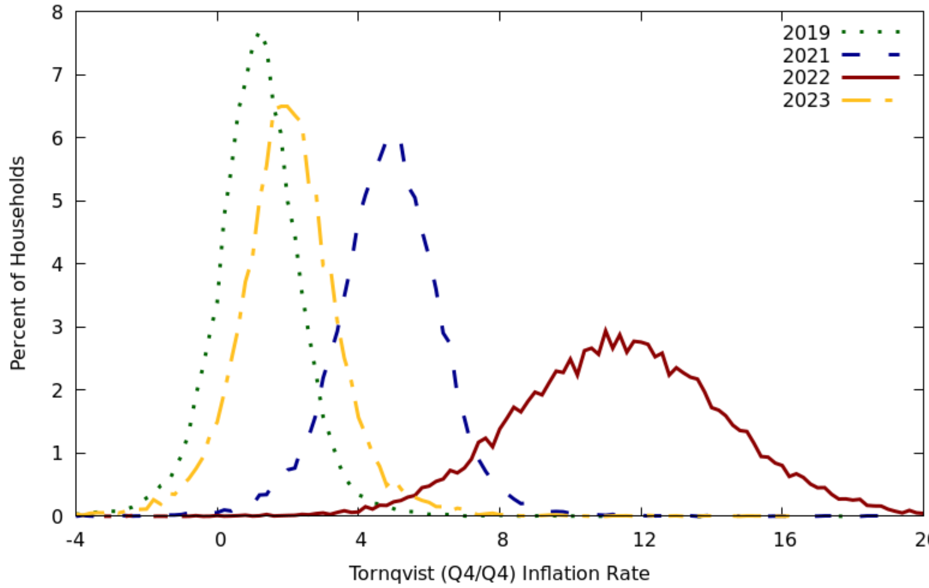

Household inflation rates computed using individual expenditure shares and applying national average price changes for each good. Tornqvist inflation rates apply average expenditure shares in Q and Q-4. Measure includes spending in all NielsenIQ HomeScan categories except general merchandise. Distributions apply households' NielsenIQ projection factors, which weight sample to be nationally representative demographically.

Figure 5
Evolution of Grocery Inflation

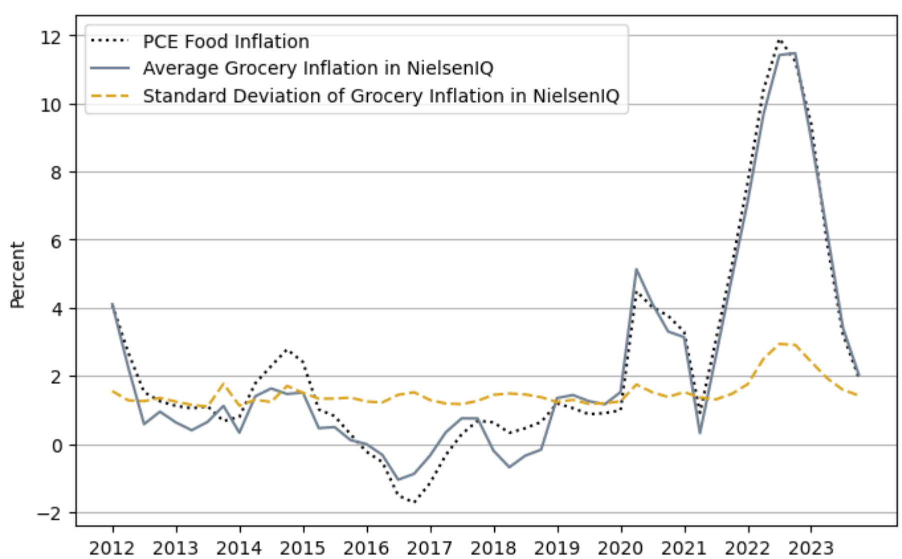

Note: Mean and standard deviation of four-quarter Tornqvist grocery inflation computed across households in NielsenIQ HomeScan panel using national average price changes for all goods (excluding general merchandise). Source: NielsenIQ and U.S. Bureau of Economic Analysis (PCE Food).

Figure 6
Distribution of Household Grocery Inflation by Demographic Group: 2022q4-2023q4

(a) Income

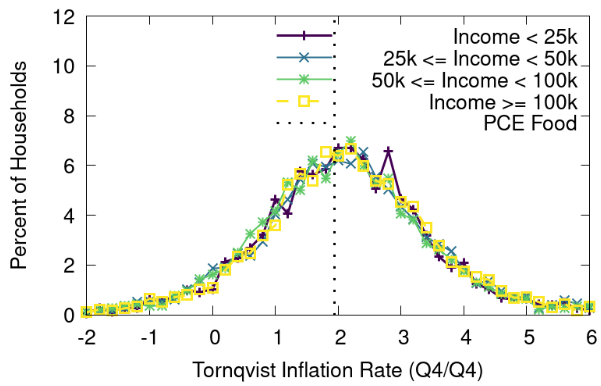

(b) Age

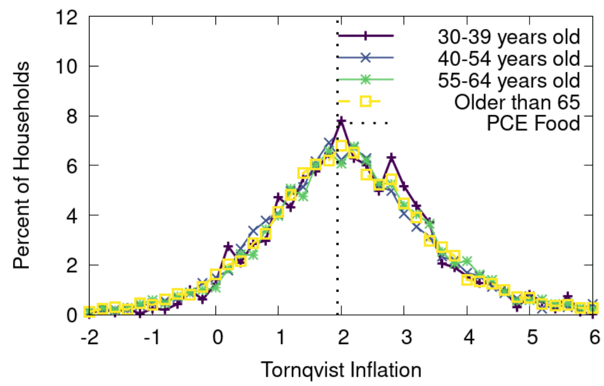

(c) Employment

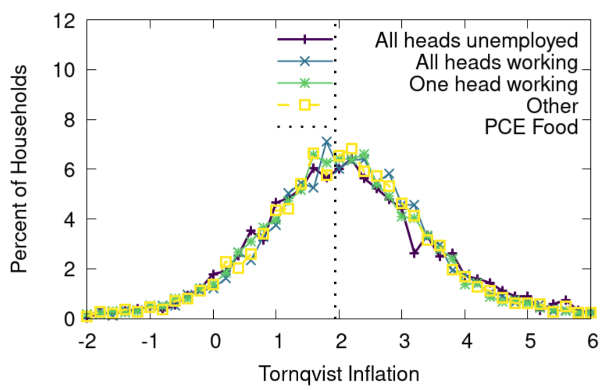

(d) Race

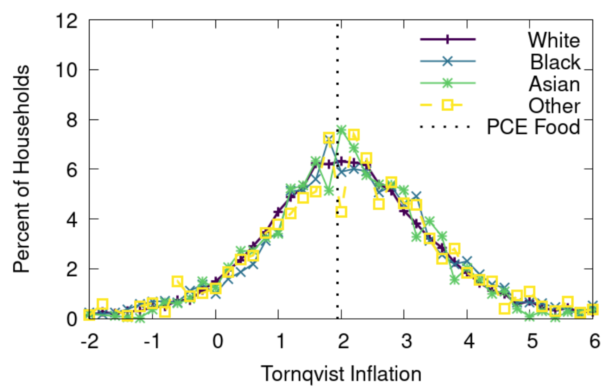

(e) Household Size

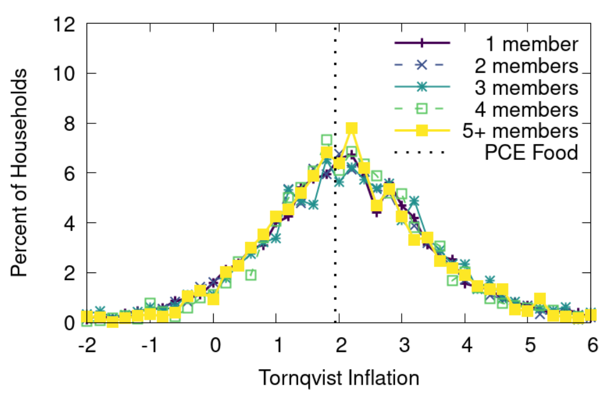

(f) ZIP Code Population Density

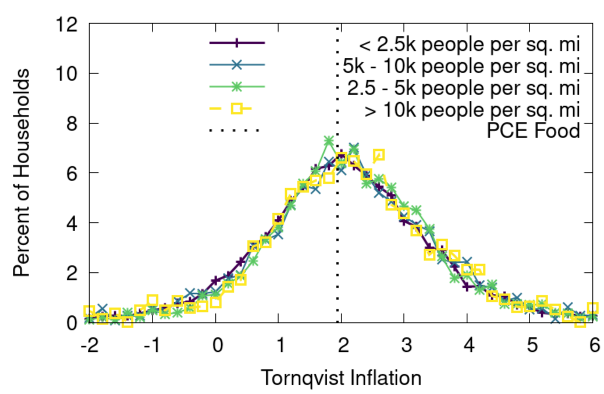

Figure 7
Discrete Choice Preferences Within Product Category

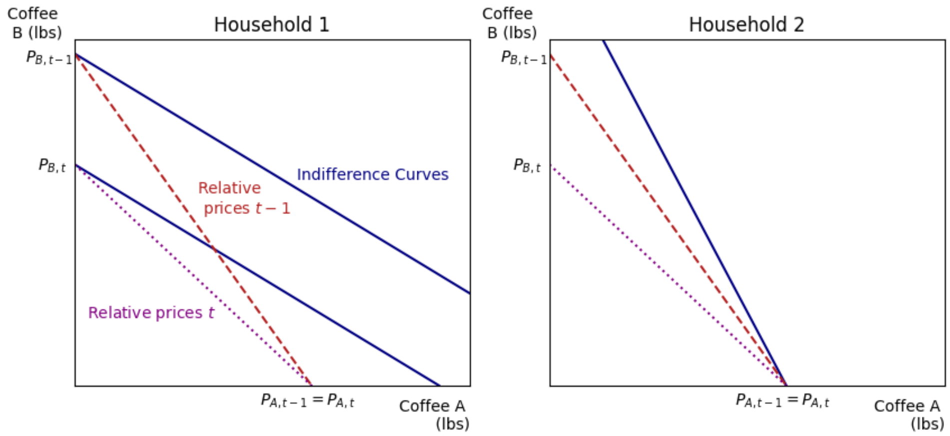

Figure shows two agents' utilities for products a and b in module m (here coffee). Navy lines represent indifference curves, whose slopes are given by $-e^{(\varphi_{ma}+\varepsilon_{hma})}/e^{(\varphi_{mb}+\varepsilon_{hmb})}$. Household 1 strongly prefers Coffee B, which received a price increase (red dashed to purple dotted budget lines). Household 2 prefers Coffee A and was unaffected.

Figure 8
Cobb-Douglas Preferences Across Product Group

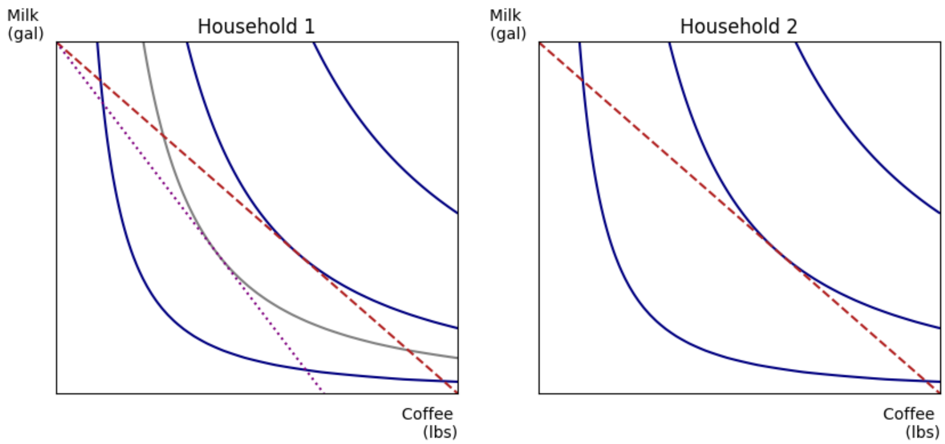

Figure shows two agents' identical Cobb-Douglas upper utility tiers. Agent 1 strongly prefers a coffee brand which received a price increase and continues to purchase it, while Agent 2 prefers a good which received no price change.

Figure 9

Relative Variance of Cumulative Cost of Living Indices

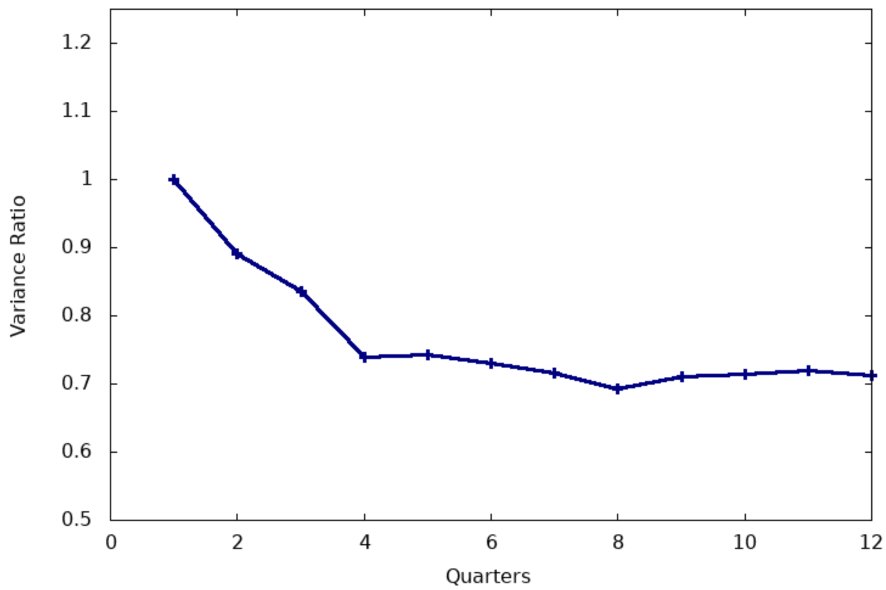

Figure presents the variance ratio (see equation 34) over horizons ranging from 1 to 12 quarters. Formally, the variance ratio at a horizon of k quarters expresses the cross-sectional variance of household inflation indices cumulating inflation to k quarters relative to k times the one-quarter variance of household inflation rates. A random walk process would exhibit a variance ratio of 1 at all horizons. A process with perfect mean reversion would exhibit a variance ratio of 0.5 for $k = 2$.

Figure 10

Distribution of Changes to Expected Lifetime Utility of Total Consumption $\mathcal{V}(\{\mathcal{C}_{h\tau}\}_{\tau=t}^{T})$

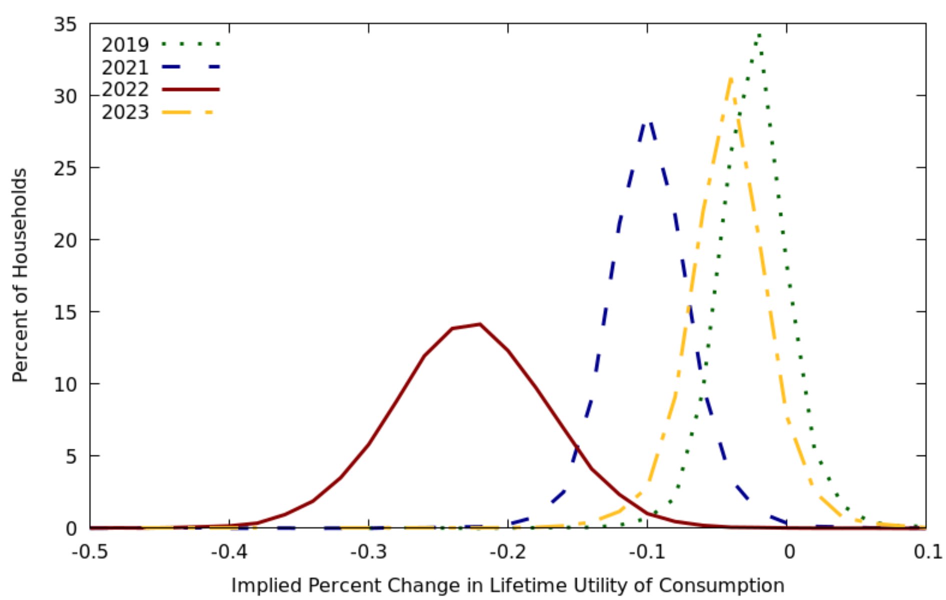

Each line translates the empirical distribution of household Q4/Q4 Tornqvist inflation experiences into changes in lifetime utility of consumption (across both grocery and non-grocery categories). Estimates assume fixed lifetime resources, and an elasticity of intertemporal substitution $\theta = 0.25$. Results assume households' costs of living take linear paths back to the average cost of living over three years.

Table 1: NielsenIQ Observations

	2017	2019	2020	2021	2022	2023
HomeScan Panel Households	62,831	61,483	60,101	58,226	56,795	55,733
... Reporting Magnet Data	52,967	53,762	52,743	50,912	49,509	48,424
... in Sample	43,693	45,010	44,518	44,162	43,463	42,445
Retail Scanner Outlets	34,134	48,903	47,532	54,093	53,620	52,459
Consumer Panel Outlets	49,507	47,976	47,400	46,559	46,353	46,498
Matched Outlets	18,828	19,995	19,577	25,010	24,271	23,938

The sample for each quarter includes households in the Consumer Panel who 1) are asked by NielsenIQ to record Magnet data (which includes categories without standardized barcodes like fresh fruits and vegetables) and 2) record at least $150 of spending in grocery categories (which exclude items in the general merchandise department such as apparel or home appliances). Matched outlets refer to those appearing in both the Consumer Panel and Retail Scanner datasets.

Table 2: NielsenIQ HomeScan Recorded Expenditures

	2019Q4	2020Q4	2021Q4	2022Q4	2023Q4
Average spending (in sample)	$1,427	$1,645	$1,644	$1,777	$1,776
Excluding general merchandise	$1,140	$1,350	$1,291	$1,418	$1,430
With any price change coverage	$1,058	$1,248	$1,171	$1,322	$1,371
With retailer price change coverage	$941	$1,104	$1,025	$1,161	$1,208
With DMA price change coverage	$661	$486	$947	$1,118	$807
On matched goods	$410	$457	$428	$461	$502

To meet sample criteria for a given quarter, households must record at least $150 in spending on non-general merchandise and purchase at least five matched goods (items with the same barcode purchased in Q and Q-4). Items only appearing in HomeScan data have retailer or DMA price change coverage if prices for that barcode are reported by at least five households and any coverage if reported by at least twenty households nationally.

Table 3: Distribution of Q4/Q4 Tornqvist Inflation by Price Change Source

Year	Percentile	Household (KSW)	Household + best alternative	Best alternative, excl. household	One price
2015	10	-8.6	-4.0	-2.2	-1.4
	25	-4.2	-1.8	-1.0	-0.6
	50	-0.6	0.2	0.0	0.2
	75	2.4	2.0	1.0	0.8
	90	6.6	4.6	2.2	1.6
2017	10	-8.4	-3.6	-1.6	-0.8
	25	-4.0	-1.4	-0.4	0.0
	50	-0.6	0.4	0.6	0.8
	75	2.8	2.6	1.8	1.6
	90	7.0	5.2	3.0	2.2
2019	10	-7.8	-3.4	-1.2	-0.4
	25	-3.4	-1.0	-0.2	0.4
	50	0.0	0.8	0.8	1.2
	75	3.4	2.8	1.8	1.8
	90	7.4	5.6	3.0	2.6
2020	10	-6.2	-2.0	0.4	1.6
	25	-1.8	0.4	1.4	2.4
	50	1.4	2.4	2.6	3.2
	75	5.0	4.8	3.8	4.0
	90	9.8	8.0	5.2	5.0
2021	10	-4.2	0.0	2.0	3.0
	25	0.2	2.6	3.2	4.0
	50	3.6	4.6	4.6	4.8
	75	7.4	6.6	5.8	5.8
	90	12.0	9.2	7.0	6.6
2022	10	0.2	5.6	7.6	7.8
	25	5.0	9.0	9.6	9.4
	50	10.8	12.0	11.6	11.4
	75	17.2	14.8	13.8	13.4
	90	25.0	17.4	15.8	15.2
2023	10	-6.2	-2.0	0.0	0.4
	25	-2.2	0.4	1.2	1.2
	50	1.2	2.2	2.2	2.0
	75	4.6	4.2	3.2	2.8
	90	8.8	6.6	4.4	3.8

Household Tornqvist inflation rates apply average expenditure shares in Q and Q-4 (and in that way incorporate goods purchased in only one period). Measure includes spending in all NielsenIQ HomeScan spending categories except general merchandise. The "One price" column conveys the paper's baseline measure depicted in Figure 4, which assumes all households pay the same national average price for an item. Other columns convey distributions under alternative price assumptions. The "Household" column uses only household price reports and therefore can be computed only on goods purchased in the base and reference quarters. The "Household + best alternative" column adds to that measure by inferring that households faced the prices at the store, retail chain, or market area where the item was purchased if such price changes are observable. In the "Best alternative, excl. household" column, this same layering takes place but households' price reports are omitted. Distributions apply households' NielsenIQ projection factors, which weight sample to be nationally representative demographically.

Table 4: Regressions of De-Meaned Four-Quarter Tornqvist Inflation

	Characteristics	+ DMA x Quarter Fixed Effects	+ Lagged Inflation
Constant	-0.091	0.079	0.004
	(0.022)	(0.079)	(0.073)
Lagged Inflation			0.431
			(0.001)
Income bin: 25-50k	-0.028	-0.031	-0.013
	(0.008)	(0.008)	(0.007)
Income bin: 50-100k	-0.096	-0.104	-0.058
	(0.008)	(0.008)	(0.007)
Income bin: $>$ 100k	-0.193	-0.205	-0.111
	(0.008)	(0.008)	(0.008)
Age bin: 30-39	0.072	0.078	0.046
	(0.021)	(0.021)	(0.019)
Age bin: 40-49	0.122	0.128	0.082
	(0.021)	(0.021)	(0.019)
Age bin: 50-64	0.146	0.153	0.099
	(0.021)	(0.020)	(0.019)
Age bin: $>$ 65	0.166	0.170	0.101
	(0.021)	(0.021)	(0.019)
Both heads unemployed	0.059	0.059	0.035
	(0.011)	(0.011)	(0.010)
Both heads working	0.028	0.027	0.022
	(0.007)	(0.007)	(0.007)
One head working	-0.002	0.002	0.003
	(0.007)	(0.007)	(0.006)
Househod size: 2	0.032	0.028	0.009
	(0.009)	(0.009)	(0.008)
Household size: 3	0.047	0.040	0.007
	(0.010)	(0.010)	(0.009)
Household size: 4	0.088	0.078	0.021
	(0.011)	(0.011)	(0.010)
Household size: 5+	0.098	0.093	0.037
	(0.011)	(0.011)	(0.010)
ZIP code density: 2.5-5k people/mile^2	-0.016	0.001	0.000
	(0.006)	(0.006)	(0.006)
ZIP code density: 5-10k people/mile^2	0.029	0.044	0.023
	(0.007)	(0.008)	(0.007)
ZIP code density: $>$ 10k people/mile^2	0.085	0.038	0.029
	(0.010)	(0.011)	(0.010)
Asian	-0.132	-0.104	-0.068
	(0.012)	(0.012)	(0.011)
Black	0.021	0.018	0.017
	(0.007)	(0.008)	(0.007)
Other non-white	-0.015	-0.008	0.003
	(0.009)	(0.009)	(0.008)
Observations	649253	649253	645885
Adjusted R^2	0.001	0.011	0.069

Note: Household inflation rates apply national average price changes and are computed among all NielsenIQ product categories except for those in general merchandise. Sample covers 2012-2023 (except 2020Q1-Q2 and 2021Q1-Q2) and includes all households for a given quarter who spend at least $150 on those categories in Q, Q-1, and Q-4. Source: NielsenIQ and U.S. Census Bureau (ZIP code densities).

Table 5: Contributions of Basket and Substitution Force to Inflation Heterogeneity

	Var(ln(Ψ)) =	Var(Basket) +	Var(Substitution) +	2 Cov(Basket, Substitution)
Value	0.000545	0.000527	0.000243	-0.000225
Percent of Total	100	97	45	-41

Table displays a decomposition of cross-sectional variance in four-quarter inflation $\Psi_{h,t-4,t}$ into contributions from basket effects (Laspeyres indices) and substitution effects (differences between Tornqvist and Laspeyres indices).

Table 6: Idiosyncratic Values of Favorite Goods Implied by CES Elasticity σ_m

Percentile of σ_m among modules	σ_m	Percentile of $\max(\varepsilon_{mhk})$ among households			
		50th %ile	75th %ile	90th %ile	95th %ile
5th	5	104%	135%	180%	215%
25th	10	48%	59%	75%	87%
50th	20	23%	28%	34%	39%
75th	40	11%	13%	16%	18%
95th	120	4%	4%	5%	6%

Values represent $100[max(e^{\varepsilon_{mhk}}) - 1]$ where ε_{mhk} are simulated for 10,000 households making 50 draws each from a Type I extreme value distribution with mean 0 and shape parameter $^1/_{(\sigma_m\, -\, 1)}$. Supposing elasticity of demand $\sigma_m = 20$, a 90th percentile value of 34% means that 10% of households value a unit of their favorite good in this product group at least 34% more than the average good of the same quality.

Table 7: Distribution of Annualized Grocery Budget Shocks

Percentile	10	25	50	75	90
2023	$31	$94	$157	$236	$299
2022	$573	$705	$852	$998	$1145
2021	$212	$283	$340	$411	$467
2020	$107	$161	$214	$268	$334
2019	$-26	$26	$77	$116	$168

Values represent annualized grocery budget shocks for households with average Consumer Expenditure Survey grocery spending in the base period and a Q4/Q4 Tornqvist grocery inflation rate at the given percentile. Grocery inflation computed using national average price changes among all NielsenIQ categories except general merchandise.

Table 8: Lifetime Utility Loss from 4% Uncompensated Grocery Inflation

θ	Permanent –	10 Years δ = -0.02	5 Years δ = -0.05	3 Years δ = -0.08	2 Years δ = -0.12	1 Years δ = -0.24	1 Quarter δ = -0.95
0.50	-0.49%	-0.22%	-0.13%	-0.09%	-0.07%	-0.04%	-0.02%
0.25	-0.49%	-0.22%	-0.15%	-0.10%	-0.08%	-0.06%	-0.04%
0.17	-0.49%	-0.23%	-0.16%	-0.12%	-0.09%	-0.07%	-0.05%

Each value conveys, for the values of risk tolerance θ given in its row and the persistence of inflation shocks given in its column, the percent change in lifetime utility from a 4% grocery idiosyncratic inflation shock under constant expenditures and zero inflation in non-grocery categories. Assumed cost of living paths exhibit aggressive mean reversion, with idiosyncratic inflation rates running below average by a constant amount until the gap is closed. δ is the inflation gap needed to return to an average cost of living over the specified horizon.

Table 9: Autoregressive/Moving Average Models of Idiosyncratic Inflation

	MA(1)	AR(1)	AR(2)	AR(3)	AR(4)
Lagged idiosyncratic inflation $\tilde{\Psi}_{h,t-2,t-1} = \epsilon_{h,t-1,t}$	-0.088	-0.093	-0.066	-0.080	-0.078
		(0.001)	(0.001)	(0.001)	(0.001)
2nd lagged idiosyncratic inflation $\tilde{\Psi}_{h,t-3,t-2}$			-0.043	-0.027	-0.030
			(0.001)	(0.001)	(0.001)
3rd lagged idiosyncratic inflation $\tilde{\Psi}_{h,t-4,t-3}$				-0.084	-0.068
				(0.001)	(0.001)
4th lagged idiosyncratic inflation $\tilde{\Psi}_{h,t-5,t-4}$					0.112
					(0.001)
Observations	1305202	1305202	1043545	811438	660306
Adjusted R^2		0.009	0.005	0.012	0.024

Model estimated using one-quarter household inflation rates computed applying national average price changes using all NielsenIQ spending categories except general merchandise. Estimation window includes 2012-2023 (excluding 2020Q1-Q2 and 2021Q1-Q2). Included households must be in sample continuously for at least 3 years.

Table 10: Income Group Coefficients: Household Outcomes and Idiosyncratic Inflation

	$ln(X_t/X_{t-4})$	$ln(C_t/C_{t-4})$	# Stores	# Trips	%Δ$Coupons
Hh. infl. x (Inc. < \$25k)	0.57	-0.43	-0.01	0.06	-0.01
	(0.013)	(0.013)	(0.002)	(0.005)	(0.000)
Hh. infl. x (Inc. \$25-50k)	0.58	-0.42	0.02	0.23	-0.01
	(0.007)	(0.007)	(0.001)	(0.003)	(0.000)
Hh. infl. x (Inc. \$50-100k)	0.65	-0.35	0.02	0.12	-0.01
	(0.015)	(0.015)	(0.002)	(0.006)	(0.000)
Hh. infl. x (Inc. > \$100k)	0.50	-0.50	-0.03	-0.04	-0.01
	(0.007)	(0.007)	(0.001)	(0.003)	(0.000)

Each column represents the dependent variable of a regression on household characteristics (the same as in Table 4) and lagged inflation. Households' demeaned four-quarter Tornqvist inflation rates are interacted with income group indicators in each regression. Rows display the coefficients (and standard errors in parentheses) of each interaction separately. The sample covers 2012-2023 (except 2020Q1-Q2 and 2021Q1-Q2) and includes all households for a given quarter who spend at least $150 on those categories in Q, Q-1, and Q-4.

Supplemental Appendix

Table A.1: NielsenIQ Coverage of Consumer Basket in 2019

Category	NielsenIQ	2019 Spending	Percent of Total
Housing		20,678.99	32.80
...Shelter	X	12,190.19	19.34
...Utilities, fuels, and public services	X	4,055.08	6.43
...Household operations	X	1,569.80	2.49
...Housekeeping supplies	✓	765.56	1.21
...Household furnishings and equipment		2,098.36	3.33
......Textiles, furniture, and floor coverings	X	677.12	1.07
......Major appliances	X	321.65	0.51
......Housewares and small appliances	*	118.69	0.19
......Miscellaneous household equipment	X	980.90	1.56
Transportation		10,742.38	17.04
...Vehicle purchases	X	4,393.70	6.97
...Motor fuel and oil	X	2,094.12	3.32
...Other vehicle expenses	X	3,474.01	5.51
...Public and other transportation	X	780.55	1.24
Food		8,169.18	12.96
...Food at home	✓	4,643.31	7.37
...Food away from home	X	3,525.87	5.59
Alcoholic beverages		579.22	0.92
...Alcoholic beverages at home	✓	316.23	0.50
...Alcoholic beverages away from home	X	262.99	0.42
Tobacco and smoking products		319.99	0.51
...Cigarettes and other tobacco products	✓	308.10	0.49
...Smoking accessories and non-tobacco products	X	11.89	0.02
Healthcare		5,193.11	8.24
...Health insurance	X	3,529.35	5.60
...Medical services	X	984.12	1.56
...Drugs and medical supplies	X	679.64	1.08
Entertainment		3,089.90	4.90
...Fees and admissions	X	879.86	1.40
...Audio and visual equipment and services	X	1,000.05	1.59
...Pets and pet care		680.96	1.08
......Pet food	✓	236.26	0.37
......Other pet products and services	X	444.70	0.71
...Toys, hobbies, and playground equipment	X	140.14	0.22
...Other entertainment supplies, equipment, and...	X	388.89	0.62
Personal care		786.40	1.25
...Personal care products	✓	418.75	0.66
...Personal care services	X	367.65	0.58
Apparel and related services		1,882.96	2.99
...Clothes for men, women, and children	X	1,227.06	1.95
...Footwear	X	418.74	0.66
...Jewelry, laundry and repair services	X	237.16	0.38
Education		1,443.25	2.29
...Tuition, loans, and tutoring	X	1,333.08	2.11
...College books and supplies	X	39.48	0.06
...Non-college books and supplies	*	70.69	0.11
Reading	X	92.08	0.15
Personal insurance and pensions		7,164.73	11.37
...Life and other personal insurance	X	519.55	0.82
...Deductions for Social Security	X	5,240.52	8.31
...Other payroll deductions	X	672.52	1.07
...Non-payroll deposit to retirement plans	X	732.14	1.16
Cash contributions		1,995.24	3.17
...To children	X	360.29	0.57
...To charitable, religious, and educational or...	X	1,161.51	1.84
...To political organizations	X	27.49	0.04
...Other cash gifts	X	445.95	0.71
Miscellaneous fees and charges	X	898.98	1.43
Total	6,877.59	63,036.41	100

Key: X = not covered, ✓= included, * = covered but not included
Source: U.S. Bureau of Labor Statistics.

A. Data

A.1 Data Processing Details

Sample Details

The most binding requirements for inclusion in the analysis sample include reporting Magnet data—products lacking barcodes—and reporting a minimum amount of spending. For households to be included in the sample for a given quarter, they must spend at least \$150 in that quarter, the previous quarter, and four quarters prior on categories other than general merchandise.

Following Kaplan and Schulhofer-Wohl (2017), households also must purchase at least five of the same goods in the base and reference quarters. Four-quarter inflation rates for a given household exclude any goods whose reported prices change by a factor of three or more.

Price Change Computations

This paper's headline indices rely on national average price changes. Sourcing those price changes between the HomeScan panel and the Retail Scanner data requires balancing each dataset's advantages and disadvantages. The HomeScan panel's sample design makes the prices its households report more representative than the Retail Scanner data. While more representative, the HomeScan panel provides less exhaustive coverage of price changes, however. With so many more products available than households in the HomeScan data, some products will lack sufficiently many reports to form an accurate estimate of the average price change. Average prices in such cases come from the Retail Scanner data.

When more than 20 households in the HomeScan data record purchasing a given barcode in both the base and reference quarters, these reports inform the average price change for the given item. This data benefits from the automatic price reports of Retail Scanner outlets when possible, but also reflect purchases at non-participating retailers. As a result, the price changes should be more insulated from any selection bias in the Retail Scanner data. For similar reasons, average price changes computed using the Retail Scanner data exclude information from retailers not visited by households in the HomeScan data when shopping for goods in that product group. The price changes of specific retailers or market areas follow similar procedures, but require five or more HomeScan panelists to record prices in the base and reference quarters.

Private label (generic) goods utilize different codes in the Retail Scanner and HomeScan data. These goods' average prices are inferred using the same rules as for retailers not in the Retail Scanner data. Such cases apply the average price from the retailer of purchase in the HomeScan data so long as at least five households report prices at the given retailer. (Others are excluded from analysis, as in cases where no prices are observed in either the reference or base quarter.)

B. Other Estimates of Inflation Heterogeneity

B.1 Applying Alternative Price Changes

The results of this section drop assumptions around the law of one price and admit price heterogeneity across households. In so doing, they introduce two new measures of household inflation which overcomes the limited focus on matched goods without losing information on price heterogeneity where it is available.

One measure uses households' reported log price changes $ln(p_{h,k,t}) - ln(p_{h,k,t-j})$ for any matched goods k that are purchased in both t and $t-j$. For unmatched goods, they employ the paradigm of replacing households' log price change with $\Delta ln(\widehat{p_{h,k,t-j,t}})$, the log price change a household *would have faced* had it bought good k in each period.[39] Wherever possible, this change should be taken from the most relevant data source available. Because the NielsenIQ Consumer Panel records the store where an item is purchased, and because many of these same stores are covered by the NielsenIQ Retail Scanner data, it is often possible to compute $\Delta ln(\widehat{p_{h,k,t-j,t}})$ as the average price change at the specific store or retailer where household h buys coffee. Incorporating such average price changes boosts coverage to 85% of expenditures in categories considered for 2019.[40]

Figures A.1 - A.3 below chart the distribution of this measure layering in several successive price changes: only household price reports, household and retailer reports, or the best prices available (household, retailer, DMA, or national average). It presents two other measures of household inflation differences as well. One, shown in yellow, applies only the average prices observed in that household's market area. The series has less variability than those applying individual prices, but more variability than the national average prices.[41] A final measure, shown in khaki, applies average price changes observed in the household's store, retailer, or market area as available before reverting to the national price. Unlike in the "best prices," measure, it skips households' specific price reports in that exercise. The result is an inflation measure which captures differences between where households live or shop where possible, but excludes other sources of price heterogeneity such as purchase timing with respect to promotional sales.

[39] Price *changes* are imputed directly in a method of double imputation, rather than the single imputation method of mixing an observed price with an imputed price, for instance by computing $ln(\hat{p}_{h,k,t}) - ln(p_{h,k,t-j})$. This method is believed to reduce error in hedonic price indices that face similar missing price problems (Hill and Melser, 2008, p. 599-600). In the household context, double imputation prevents the accidental introduction of noise. A household always buying a good in bulk when it goes on sale will have a below-average observed price $ln(p_{h,k,t-j})$. Comparing this directly to the average price $ln(\hat{p}_{h,k,t})$ likely overstates the household's effective price change. The average price change $ln(\widehat{p_{h,k,t-j,t}})$ will tend to understate inflation heterogeneity, but does not accidentally introduce price change volatility.

[40] A subtle complication arises if household h bought Almond Breeze at one store and Berkeley Farms 2% milk at another. In such instances, take the average price change of Almond Breeze from the first store and the average price change of Berkeley Farms 2% milk from the other.

[41] It is possible, with such detailed data, that variability in DMA average prices of some items reflect small samples rather than differing price availability. To mitigate this, valid price changes must be based on the reports of at least five households in each period.

Figure A.1
Household Inflation 2018q4-2019q4 by Price Change Source

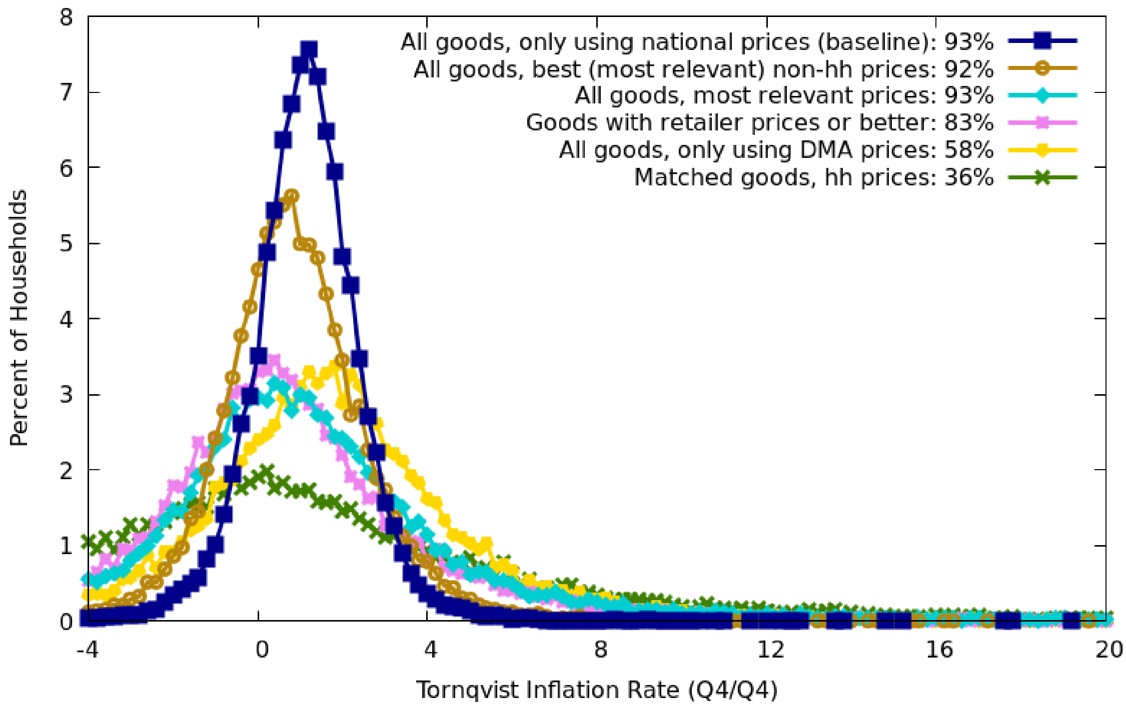

Notes for Figures A.1 - A.4: each line conveys a measure of inflation constructed using a different source of price change information. The "matched goods, hh prices" line employs only price information reported by households, as in Kaplan & Schulhofer-Wohl (2017), and as a result is limited to the set of goods purchased in both the base and reference periods. The "goods with retailer prices or better" line uses household price reports when available, but applies the average quarterly price change at the store or retailer of purchase for goods not purchased in both periods. The "all goods, best prices" measure continues to layer in the next best available price, either specific to the household's market area or the national average. In contrast to each of these measures, the "all goods, best non-hh prices" construction similarly uses the best available price for a good, but skips the household's individual price reports. As a result, it excludes idiosyncratic price differences across households (for instance arising from promotional sales and purchase timing) but retains heterogeneity driven by outlet or location. The "only using DMA prices" line applies average price changes for each item within its market area while the baseline measure assumes that all households face the same average price change for a given barcode.

Figure A.2
Household Inflation 2020q4-2021q4 by Price Change Source

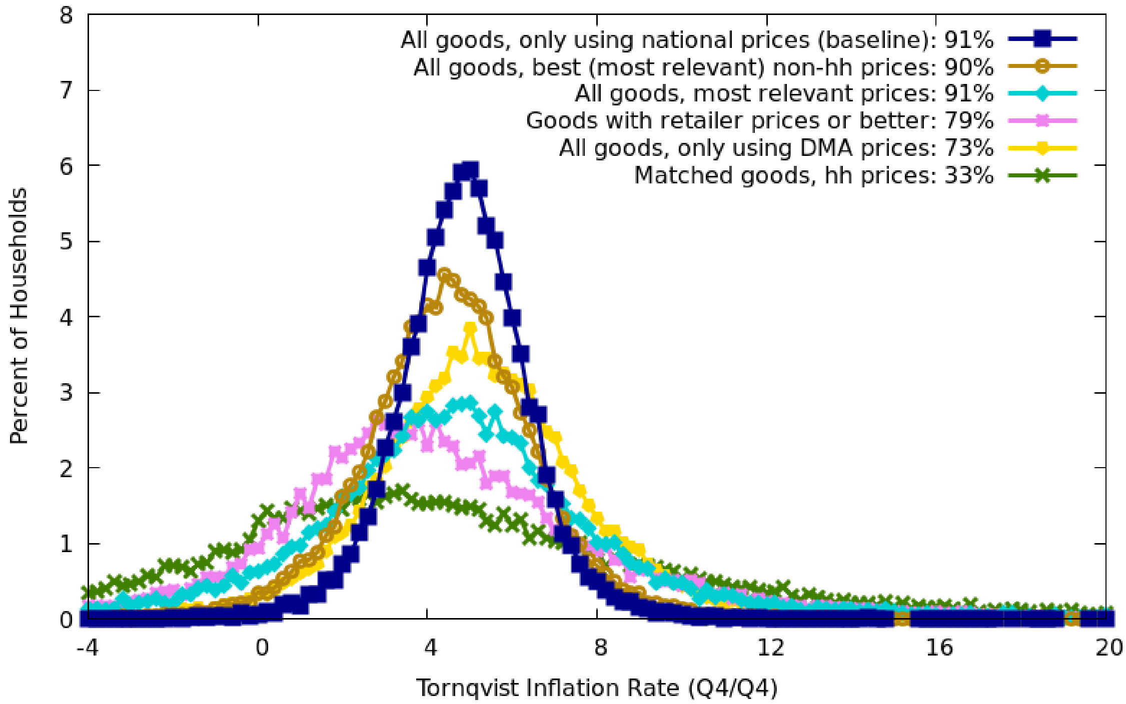

Notes for Figures A.1 - A.4: each line conveys a measure of inflation constructed using a different source of price change information. The "matched goods, hh prices" line employs only price information reported by households, as in Kaplan & Schulhofer-Wohl (2017), and as a result is limited to the set of goods purchased in both the base and reference periods. The "goods with retailer prices or better" line uses household price reports when available, but applies the average quarterly price change at the store or retailer of purchase for goods not purchased in both periods. The "all goods, best prices" measure continues to layer in the next best available price, either specific to the household's market area or the national average. In contrast to each of these measures, the "all goods, best non-hh prices" construction similarly uses the best available price for a good, but skips the household's individual price reports. As a result, it excludes idiosyncratic price differences across households (for instance arising from promotional sales and purchase timing) but retains heterogeneity driven by outlet or location. The "only using DMA prices" line applies average price changes for each item within its market area while the baseline measure assumes that all households face the same average price change for a given barcode.

Figure A.3
Household Inflation 2021q4-2022q4 by Price Change Source

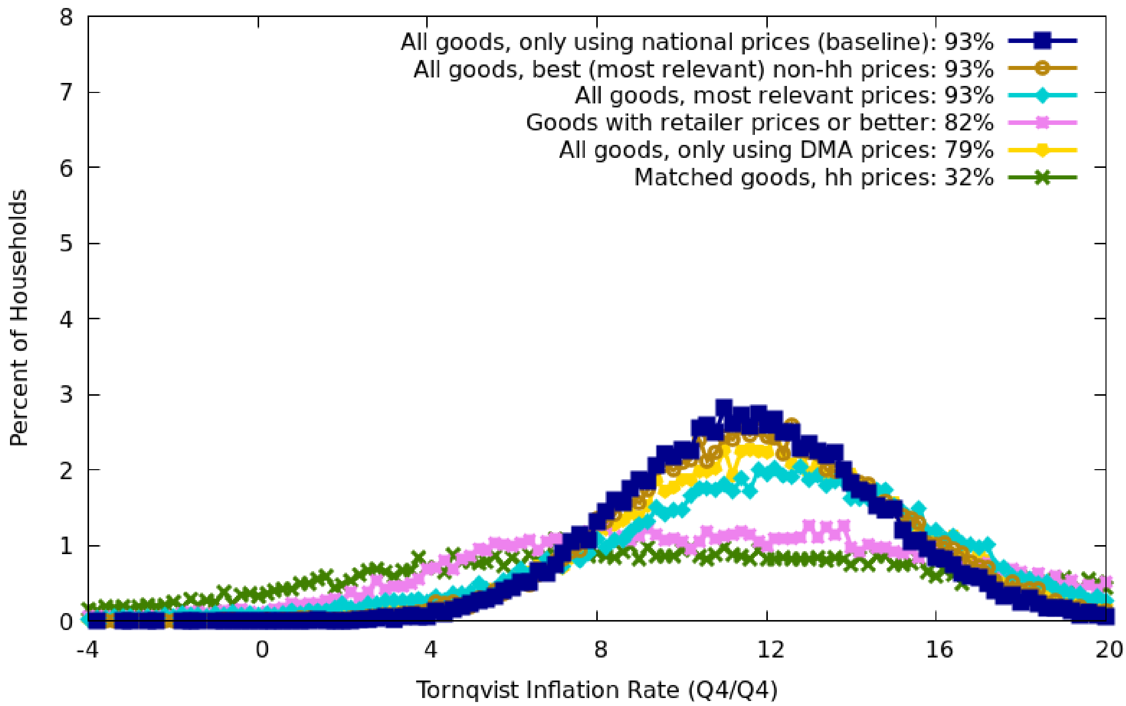

Notes for Figures A.1 - A.4: each line conveys a measure of inflation constructed using a different source of price change information. The "matched goods, hh prices" line employs only price information reported by households, as in Kaplan & Schulhofer-Wohl (2017), and as a result is limited to the set of goods purchased in both the base and reference periods. The "goods with retailer prices or better" line uses household price reports when available, but applies the average quarterly price change at the store or retailer of purchase for goods not purchased in both periods. The "all goods, best prices" measure continues to layer in the next best available price, either specific to the household's market area or the national average. In contrast to each of these measures, the "all goods, best non-hh prices" construction similarly uses the best available price for a good, but skips the household's individual price reports. As a result, it excludes idiosyncratic price differences across households (for instance arising from promotional sales and purchase timing) but retains heterogeneity driven by outlet or location. The "only using DMA prices" line applies average price changes for each item within its market area while the baseline measure assumes that all households face the same average price change for a given barcode.

Figure A.4
Household Inflation 2022q4-2023q4 by Price Change Source

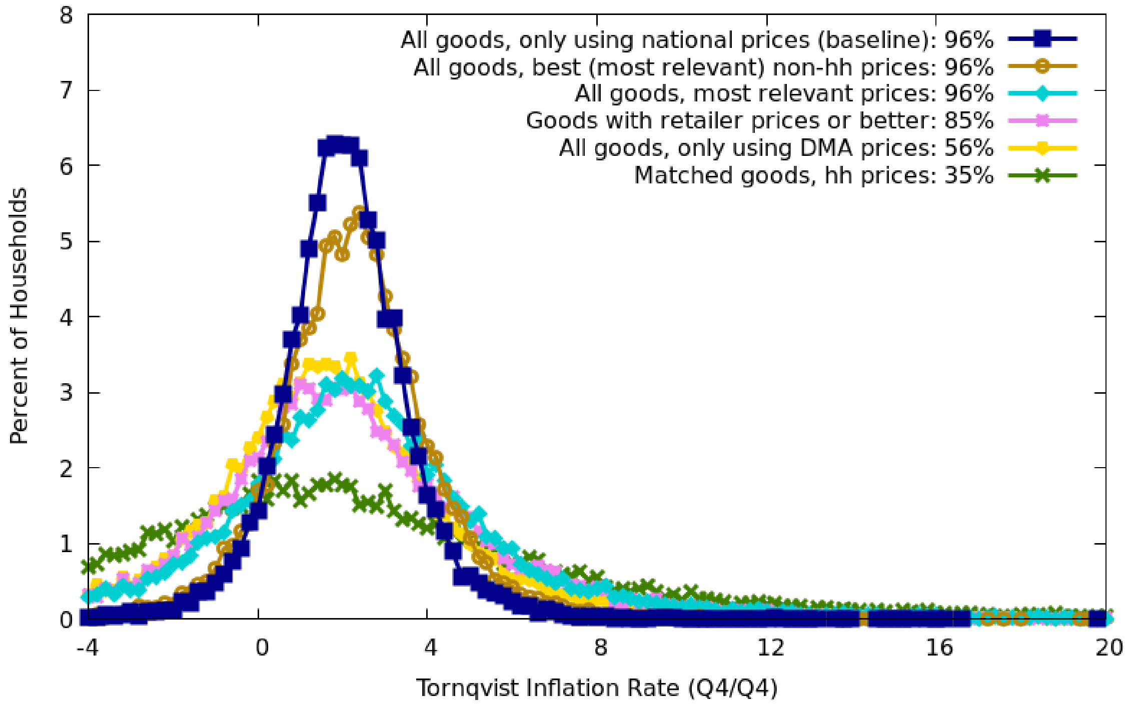

Notes for Figures A.1 - A.4: each line conveys a measure of inflation constructed using a different source of price change information. The "matched goods, hh prices" line employs only price information reported by households, as in Kaplan & Schulhofer-Wohl (2017), and as a result is limited to the set of goods purchased in both the base and reference periods. The "goods with retailer prices or better" line uses household price reports when available, but applies the average quarterly price change at the store or retailer of purchase for goods not purchased in both periods. The "all goods, best prices" measure continues to layer in the next best available price, either specific to the household's market area or the national average. In contrast to each of these measures, the "all goods, best non-hh prices" construction similarly uses the best available price for a good, but skips the household's individual price reports. As a result, it excludes idiosyncratic price differences across households (for instance arising from promotional sales and purchase timing) but retains heterogeneity driven by outlet or location. The "only using DMA prices" line applies average price changes for each item within its market area while the baseline measure assumes that all households face the same average price change for a given barcode.

Figure A.5
Comparison of Baseline and Matched Goods Inflation

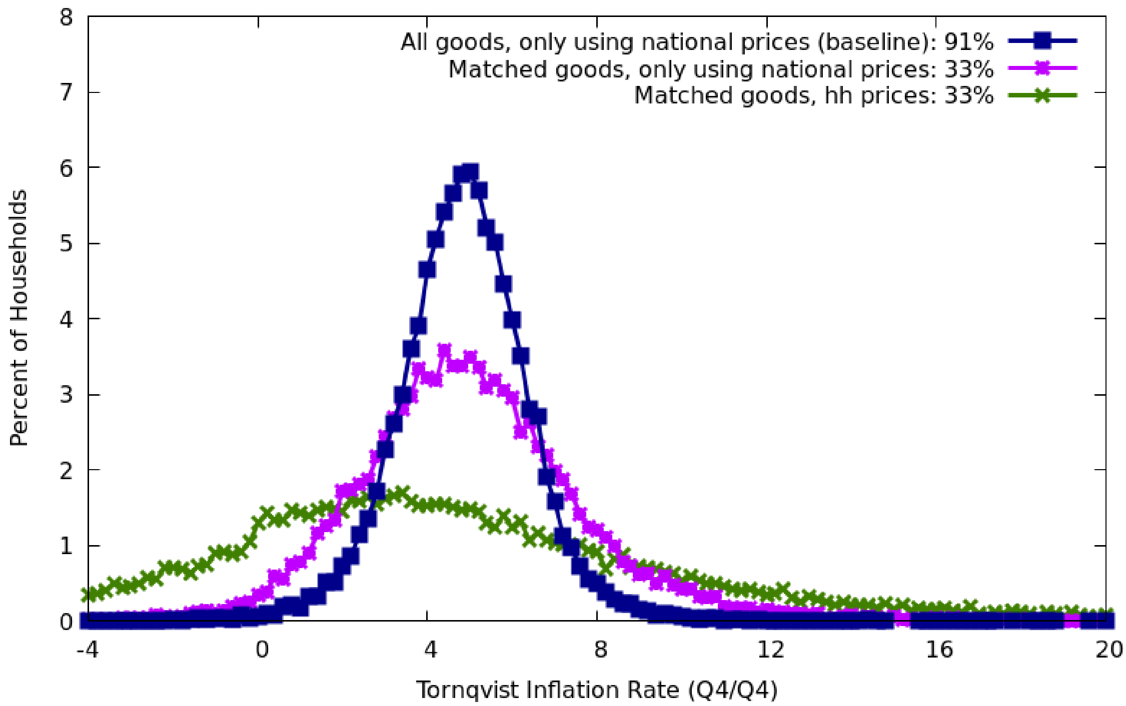

Figure plots the distribution of household Tornqivst inflation from 2020Q4-2021Q4 in a comparison of three measures. Its navy line, which assumes all households face the national average price of each barcode, represents the paper's baseline measure. It computes inflation over all goods for which national average price changes can be observed. The green line applies only household-level price changes, necessarily restricting attention to matched goods that households purchase in both 2020Q4 and 2021Q4. A magenta line decomposes the differences between these two measures by applying national average price changes (as in the navy line) to matched goods (as in the green).

B.2 Paasche vs. Laspeyres Indices

To provide a sense of how strongly household substitution patterns affect inflation rates, Appendix Figure A.6 represents a version of Figure 4 which replaces Tornqivst indices with the Laspeyres and Paasche indices. Laspeyres indices, recall, assume no benefit from realized substitution behaviors; Paasche indices assume no cost of substituting. In 2019 and 2021, substitution introduces almost no gap between these inflation distributions. A curious pattern emerges in 2022, however. Households' substitution behaviors between 2021Q4 and 2022Q4 tended to contribute positively to their overall inflation rate.

Some portion of this effect appears to be due to households substituting in the "wrong direction" more often. Figure A.7 plots the distribution of households' expenditure shares on product modules where Laspeyres inflation exceeds Paasche inflation (Lasp. > Psch.) or have inflation rates less than one percentage point apart between the measures (Lasp. - Psch. > 1%). Meaningful substitution in the wrong direction occurs in a minority of cases, but that rate rose in 2022. Modules where Paasche inflation exceeded Laspeyres inflation by more than one percentage point accounted for

Figure A.6
Distribution of Laspeyres and Paasche Inflation

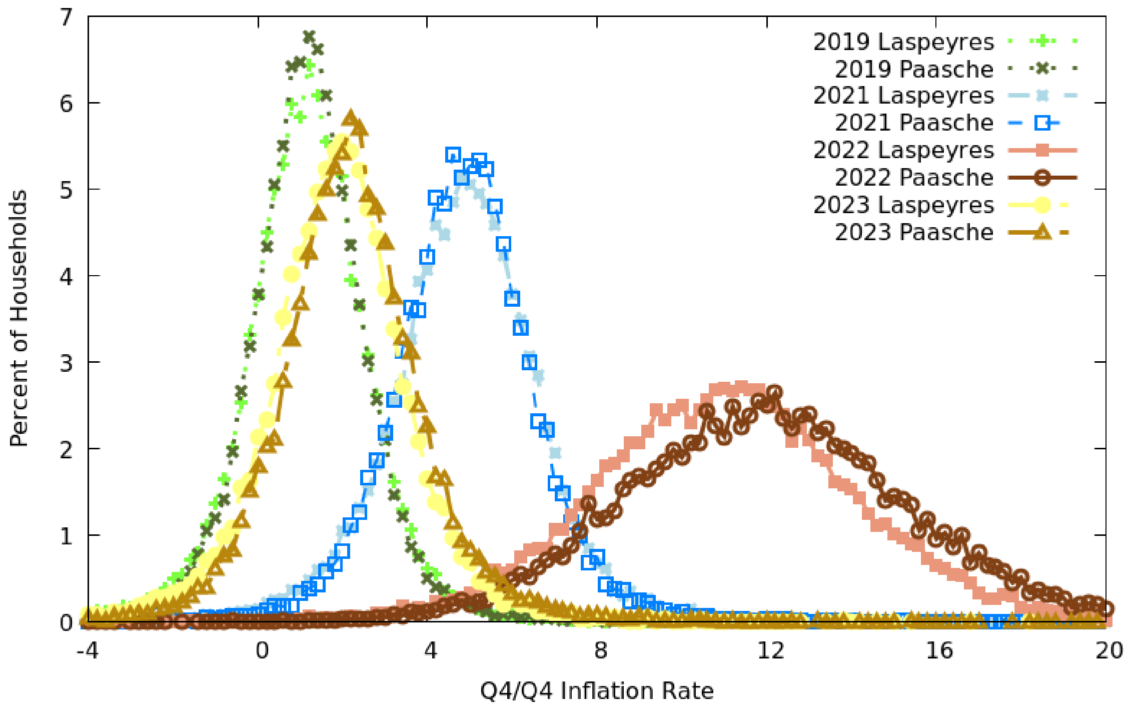

Distributions of Laspeyres and Paasche Q4/Q4 inflation rates among households in NielsenIQ HomeScan panel. Inflation computed using national average price changes for all goods except those in general merchandise categories.

22% of expenditures in 2019. In 2022, however, this figure rose to 30%. Not plotted, the share of spending for which Paasche inflation exceeds Laspeyres inflation by more than 5 percentage points rose from 8% to 18%.

When households substitute in ways that do not lower their inflation rates, it may reflect patterns of "cheapflation," a phenomenon Cavallo and Kryvtsov (2024) document using web-scraped price data. They find that from 2021-2023, goods at the low end of the quality distribution experienced significantly higher price changes than those at the high end (Fig. 5, p. 9). In such an environment, households whose preferred items move out of their budget set will sometimes be forced to switch to an alternative which exhibited an even higher price change. For this case, the Paasche index continues to bound households' inflation costs from below. Laspeyres indices, which assume the household can consume some fractional portion of their preferred goods, may fail to fully capture the loss of those goods moving out of the budget set.

Figure A.7
Distribution of Spending Shares in Modules With Laspeyres > Paasche

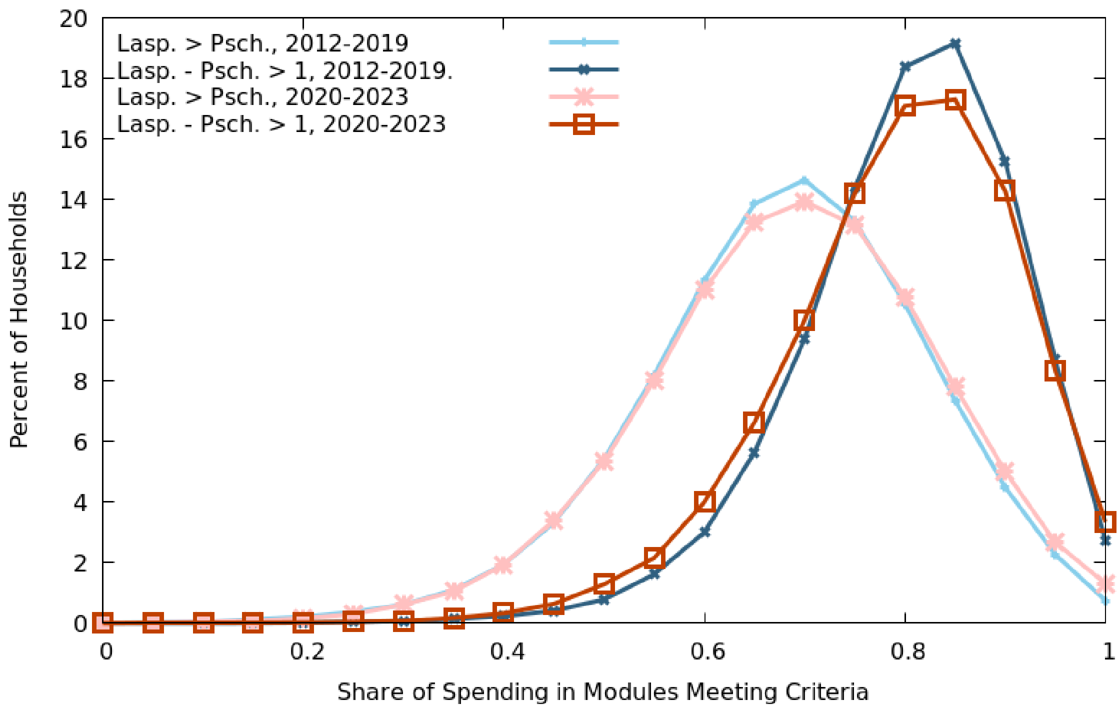

Figure constructed by evaluating, for each household and product module, whether the household's Laspeyres index in that module exceeds (Lasp. > Psch.) or is within one percentage point (Lasp. - Psch. > 1) of its Paasche index. The statistic whose distribution these lines plot are households' shares of grocery expenditures meeting that condition.

C. *Inflation Heterogeneity by Demographic Group*

Figure A.8
Distribution of Household Inflation by Demographic Group:
2020q4-2021q4

(a) Income

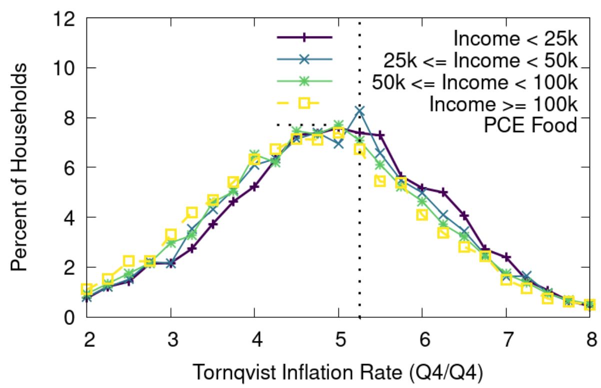

(b) Age

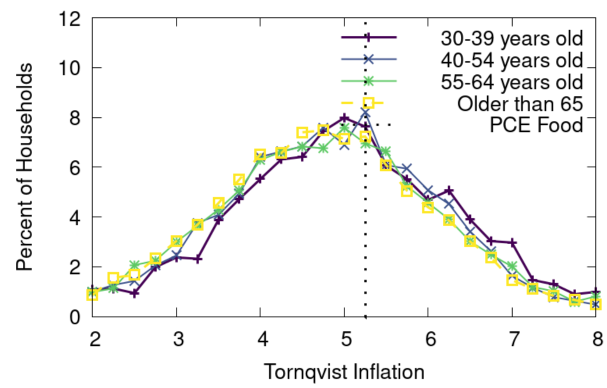

(c) Employment

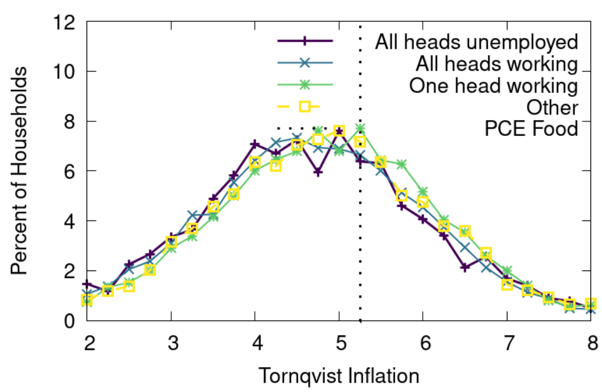

(d) Race

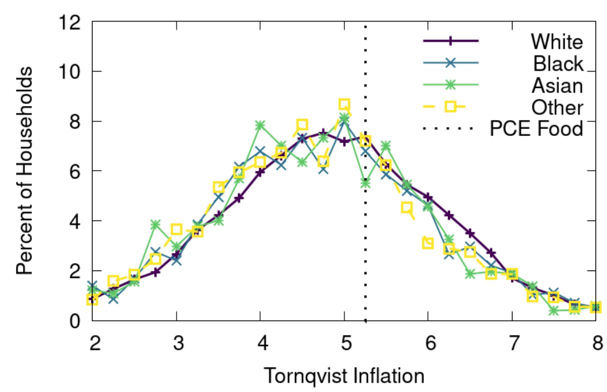

(e) Household Size

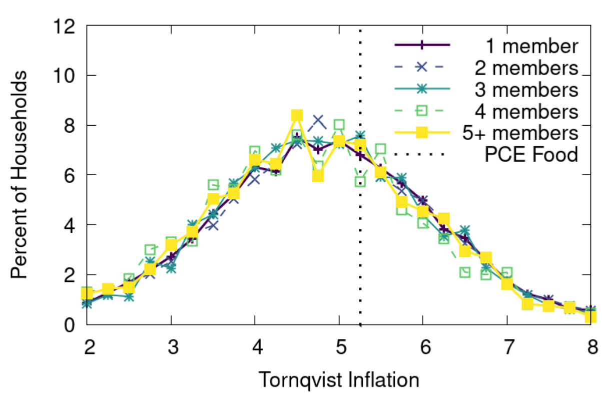

(f) Population Density

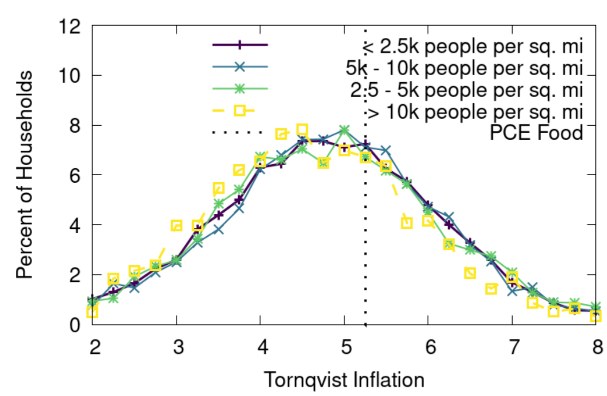

Figure A.9
Distribution of Household Inflation by Demographic Group:
2021q4-2022q4

(a) Income

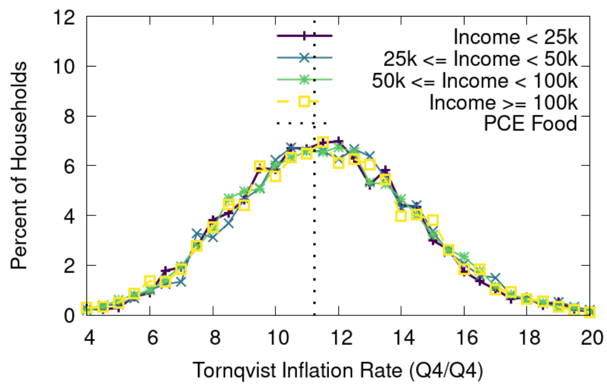

(b) Age

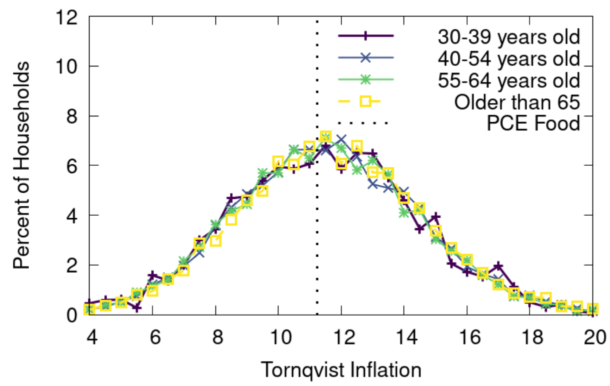

(c) Employment

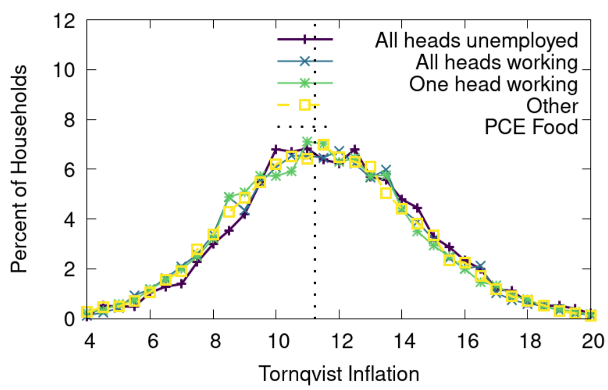

(d) Race

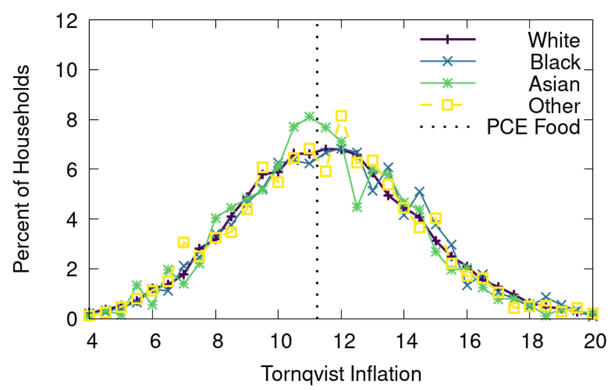

(e) Household Size

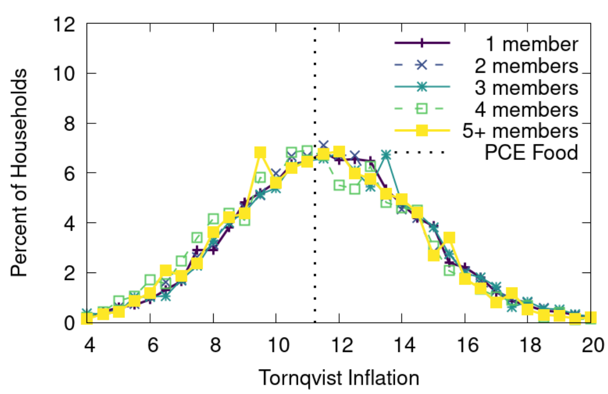

(f) Population Density

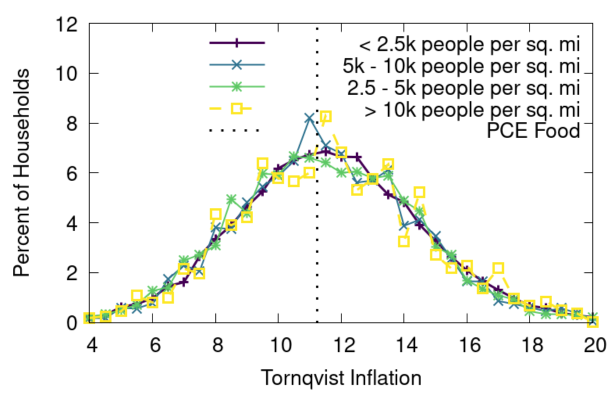

D. *Household Preferences*

D.1 Scale of Household Idiosyncratic Preferences

Table A.2: Relative Idiosyncratic Value of 1st and 2nd Most Preferred Goods

Percentile of σ_m among modules	σ_m	Percentile of value among households			
		50th %ile	75th %ile	90th %ile	95th %ile
5th	5	12%	26%	47%	65%
25th	10	7%	14%	23%	31%
50th	20	3%	7%	12%	15%
75th	40	2%	3%	6%	8%
95th	120	1%	1%	2%	3%

Values represent $100[(^{\varepsilon_{hm1}}/_{\varepsilon_{hm2}}) - 1]$ where ε_{hmk} are simulated for 10,000 households making 50 draws each from a Type I extreme value distribution with mean 0 and shape parameter $^1/_{(\sigma_m - 1)}$. Supposing $\sigma_m = 20$, a 90th percentile value of 12% means that 10% of households value a unit of their favorite good in this product group at least 12% more than their next most preferred variety of the same quality. Range of elasticities as estimated by Argente and Lee (2021).

D.2 Concavity of CRRA Preferences

Figure A.10
Total Utility of Consumption

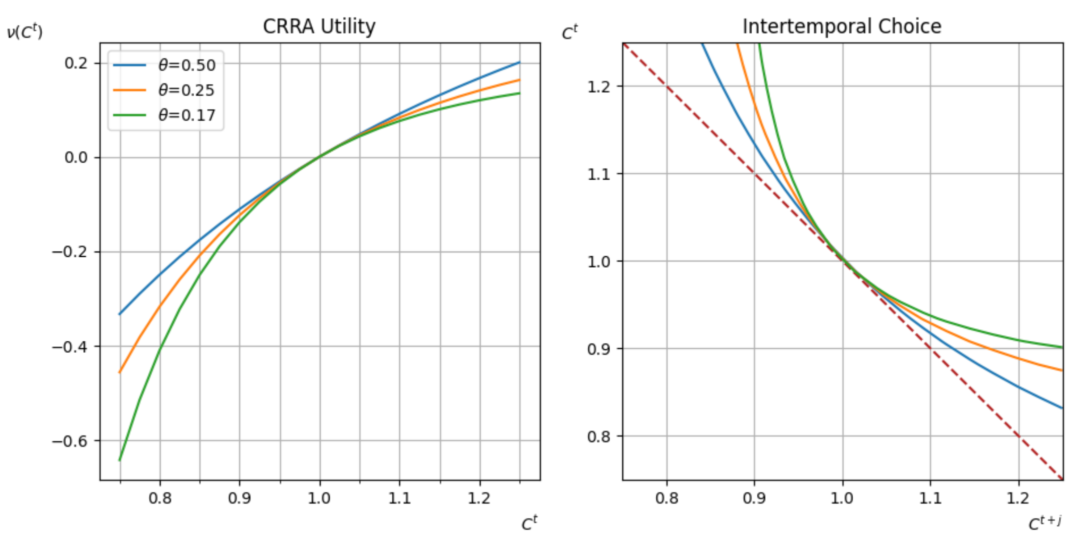

E. Costs of Inflation Heterogeneity

Figure A.11
Distribution of Changes to Flow Utility of Total Consumption $\nu(\mathcal{C}_{ht})$

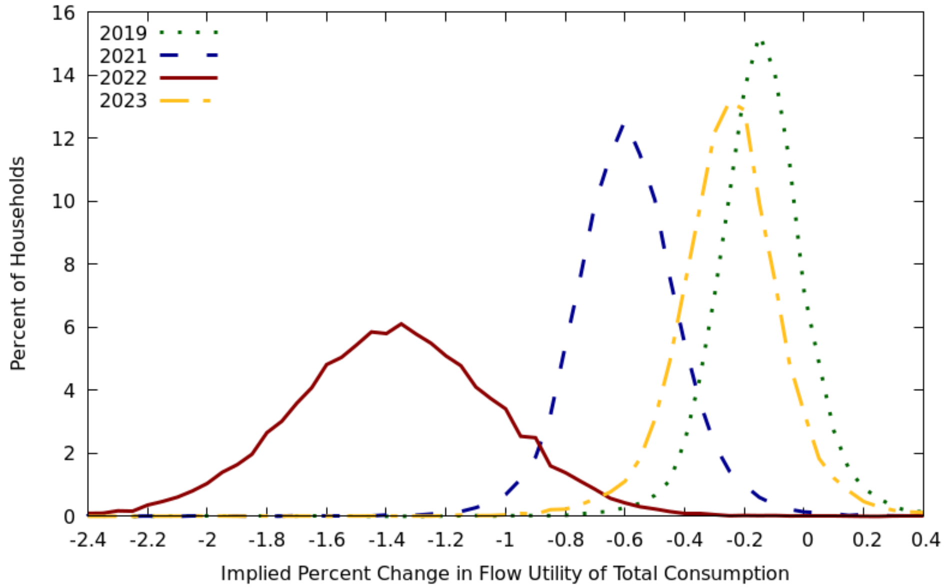

Each line translates the empirical distribution of household Q4/Q4 Tornqvist inflation experiences—as in Figure 4—into implied changes in flow utility of total consumption across both grocery and non-grocery spending categories $\nu(\mathcal{C}_{ht})$. Estimates assume a fixed budget, with expenditures not increasing to offset the price shocks, and a coefficient of relative risk tolerance $\theta = 0.25$.

E.1 Alternative Cost of Living Shock Persistence

Figure A.12
Distribution of Lifetime Utility Changes for Shocks of Varying Persistence
Absent Expenditure Increases

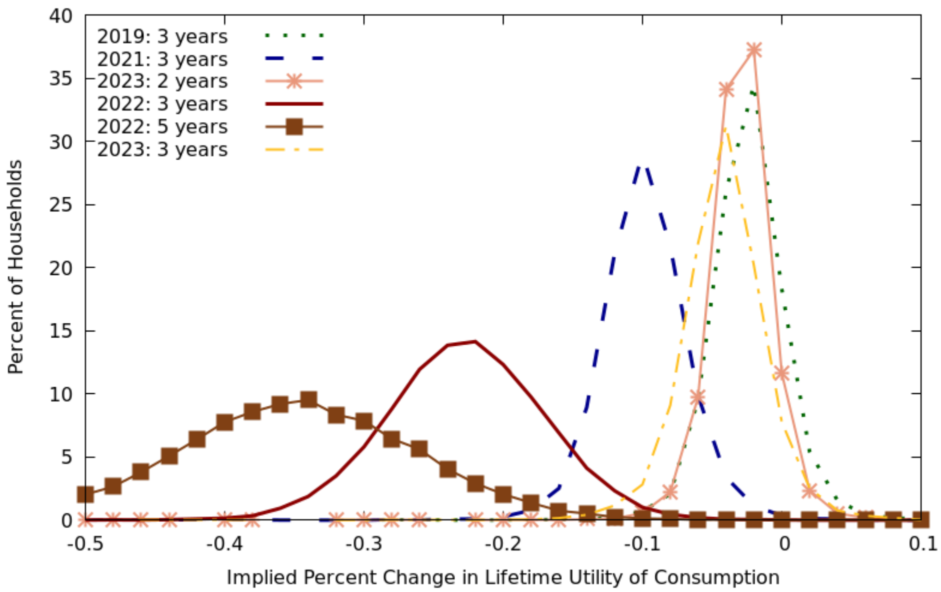

Each line translates the empirical distribution of household Q4/Q4 Tornqvist inflation experiences—as in Figure 4—into implied changes in total lifetime utility of total consumption across both grocery and non-grocery spending categories. Estimates assume fixed lifetime resources, with expenditures not increasing to offset the price shocks, and an elasticity of intertemporal substitution $\theta = 0.25$. The exercise requires an assumption about how cost of living paths respond to inflation shocks. Baseline results assume households' costs of living take linear paths back to the average cost of living over the specified duration.