ChicagoFedLetter

Chicago Fed Letter, No. 533, November 2026

Crossref

The Declining Gap in Prices Paid for Consumer Packaged Goods by Households with Higher and Lower Incomes

By Diego Jiménez Hernández, Alli Marney-Bell

Microeconomics | Macroeconomics and Monetary Economics

Historically, households with higher incomes have paid more than those with lower incomes for what economists typically call "consumer packaged goods," which are items that households purchase frequently, often coming in bags, bottles, or boxes (see, e.g., Broda et al., 2009; Blow and Leicester, 2012; and Pisano et al., 2022).[1] The gap between the prices paid for such goods by households with higher and lower incomes, however, has narrowed since 2004, with most of the convergence occurring after the first quarter of 2020. In this *Chicago Fed Letter*, we show the narrowing in this price gap has been driven almost entirely by prices becoming more similar within the same product categories for both income groups rather than by households changing their spending patterns. We show that the price gap narrowed because lower-income households faced higher inflation during this period, compared to their higher-income counterparts.

We use the NielsenIQ Consumer Panel Service—a longitudinal survey of roughly 40,000–60,000 households per month across the United States, covering around 40 million transactions for consumer packaged goods per year over the period 2004–22. For each transaction, we observe the product purchased at the barcode level, the price paid, the store type, and the household's income and demographics. This level of granularity in the data lets us compare prices paid for the same goods across households with very different incomes and track how those comparisons change over time.

We focus on the top 50 categories of packaged goods by expenditure share that we continuously observe throughout the sample period (these are listed in figure A1 of the appendix). We restrict the analysis to products measured in the most common unit within the category (e.g., fluid ounces for milk) and lightly clean the data to remove outliers.[2] This sample covers barcoded consumer packaged goods in the categories of milk, carbonated sodas, and coffee, among others.[3] Throughout the analysis, we use the sample weight provided by NielsenIQ, which accounts for the survey design and makes our results representative of all U.S. households excluding those in Alaska and Hawaii.

Price differences across the income distribution

We start by dividing the price a household pays per unit of volume (or weight) by the national average price in that product category and quarter. This yields a percentage deviation from the category mean for each transaction in the data. For example, if a household pays $1.20 for a gallon of milk of brand A in a quarter when the national average across all brands of milk is $1.00, its deviation is 20%. Since we aggregate transactions across widely different product categories with different units of volume (e.g., sodas, cookies, and toilet papers), we weight observations by their expenditure share and plot the average deviation by income group (or bin) for the first year in our data (2004) in figure 1.[4] The income bins are not equally spaced and were chosen by NielsenIQ; they run from less than or equal to $4,999 to more than or equal to $100,000 per year in nominal U.S. dollars. We also plot the price difference by income group, category by category, in figure A2 of the appendix.

1. Price difference from the national average price of the consumer packaged goods category, averaged across 50 product categories, by income group, 2004

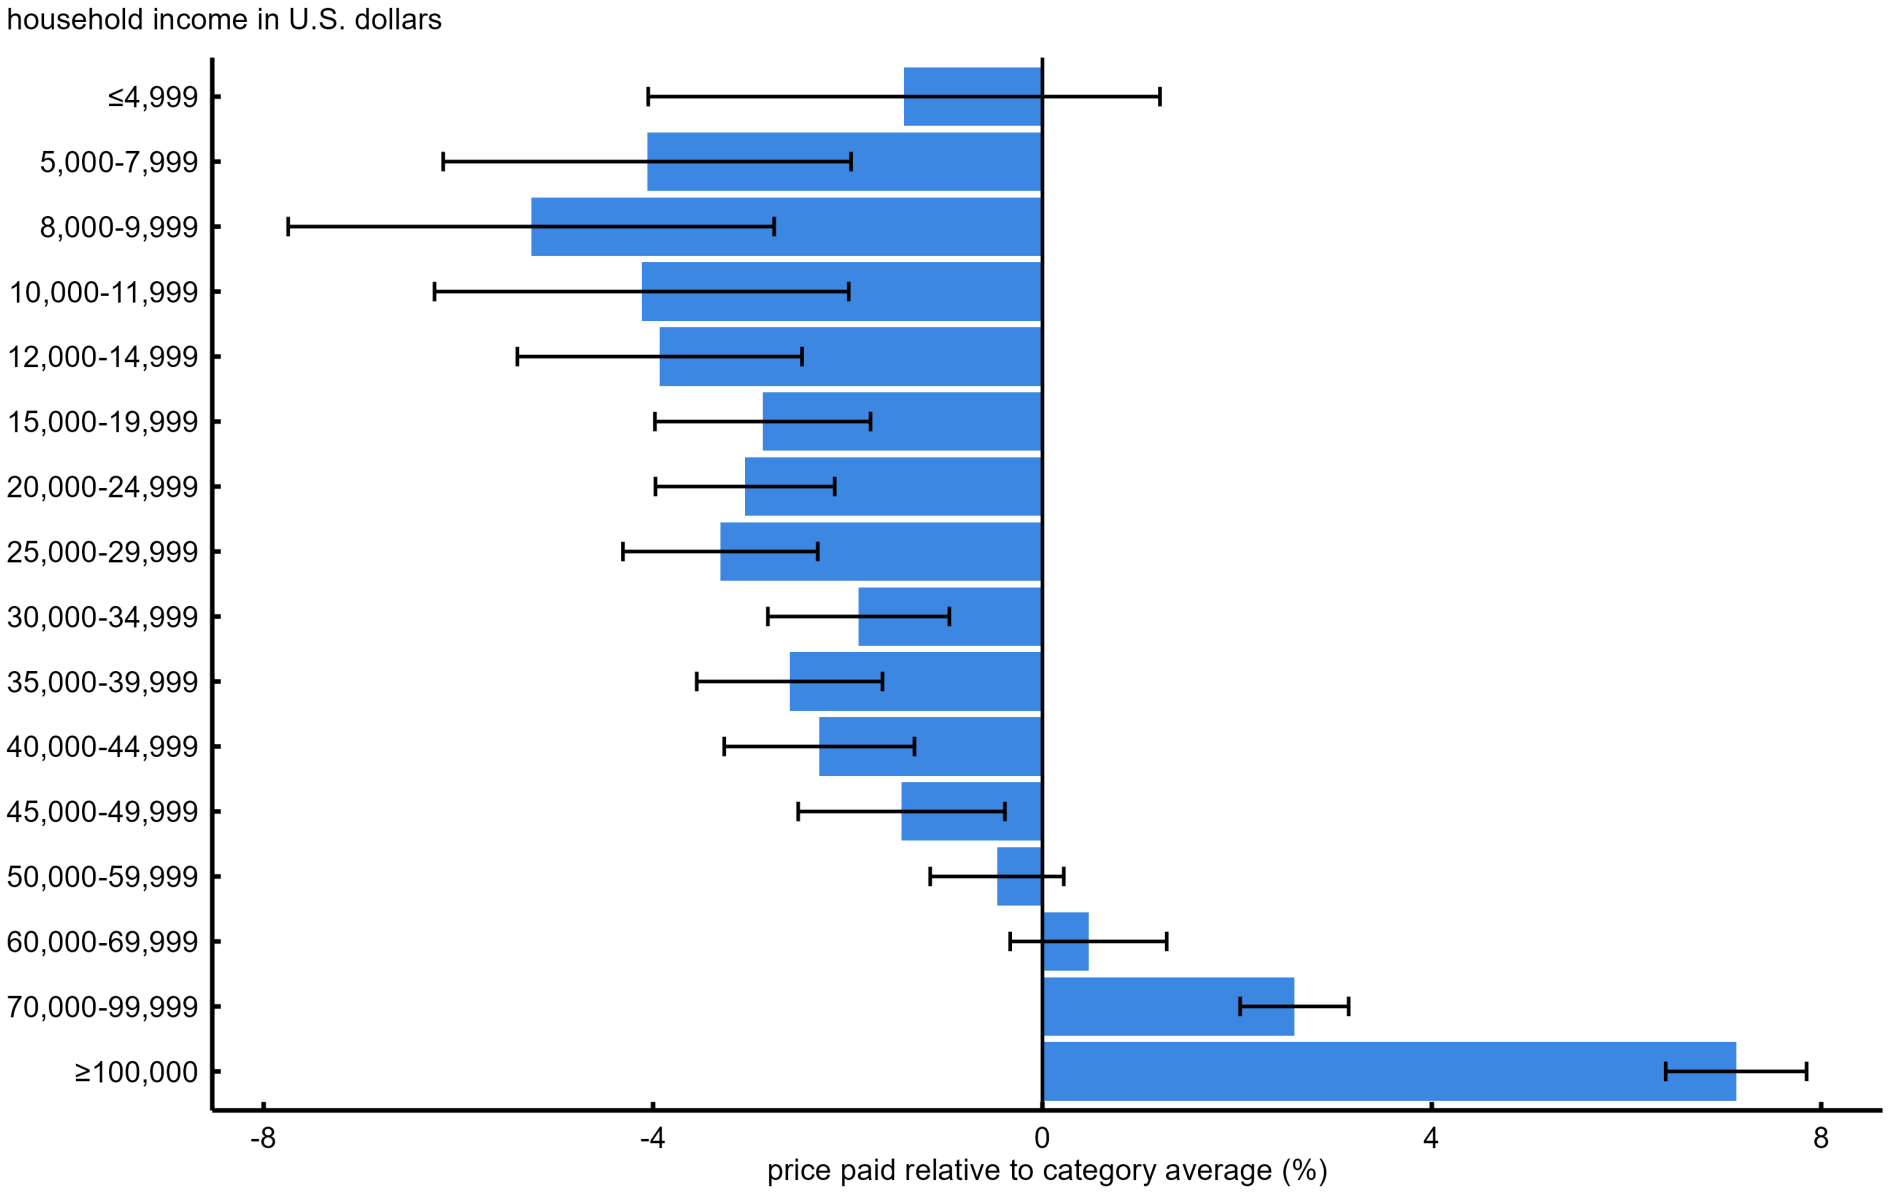

Notes: Each blue bar shows, for a given household income group (or bin), the average price it paid relative to the national average price for the same product category and quarter in 2004, expressed as a percentage and pooled across the top 50 categories by expenditure share (see figure A1 in the appendix for the complete list of 50 categories of consumer packaged goods). For example, a value of 3% means the group paid 3% more than the category average, averaged across all 50 categories. The observations are weighted by their expenditure share. The black whiskers are 95% confidence intervals based on standard errors clustered by household. The income bins are not equally spaced and follow the NielsenIQ Consumer Panel Service definitions; they run from less than or equal to $4,999 to more than or equal to $100,000 per year in nominal U.S. dollars.
Source: Authors' calculations based on data from the NielsenIQ Consumer Panel Service.

Households with higher incomes consistently pay more than average, while households with lower incomes pay less. Apart from the bottom two bins (where measurement error or partial income reporting may be an issue), the income gradient is monotonic and the percentage differences are approximately linear in income. The difference runs from about −5.3% for those earning between $8,000 and $9,999 per year to roughly 7.1% for those earning more than or equal to $100,000. The pattern is not driven by any single product category. The positive income gradient holds within the large majority of the top 50 categories rather than being concentrated in a few; figure A2 in the appendix shows the gradient is broad-based, with only a few exceptions among the top 50 categories (e.g., diet soft drinks and disposable diapers) where the gradient is mostly flat. The variation in prices relative to category averages between households with higher and lower incomes reflects many different phenomena, such as variation in the quality of products purchased within a category (e.g., a premium or organic brand rather than a store brand), quantity discounts, preference differences, and shopping intensity, to name just a few.

Differences in prices for households with higher and lower incomes over time

We now show how these price differences have evolved over time. To track them, we construct a single statistic that summarizes the price gap between households with higher and lower incomes. Take the year 2004 (figure 1) as an example. According to the *American Community Survey* (ACS), a nationally representative survey that is often used to measure the income distribution in the U.S., a household in the 75th percentile of the income distribution earned about $75,700 in 2004, while a household in the 25th percentile earned about $22,540. As shown in figure 1, the 75th percentile household's bar is roughly 5.9 percentage points higher than that of the 25th percentile household, indicating that the higher-income household paid roughly 5.9 percentage points more than the lower-income one for the same categories of goods. That price difference between the 75th and 25th percentiles of the income distribution—what we're calling the "75–25 price gap" for short—is what we want to track in a systematic way.

To construct the 75–25 price gap, we fit a linear regression model of the percentage deviation from the national category average against the logarithm of the midpoint of each income group, separately for each quarter.[5] This gives us a relationship between (log) income and percentage price differences. Since the national income quartiles are not necessarily at the center of the NielsenIQ intervals for the income bins, we use the fitted regression line to predict the price deviation (from the national category average) at the 75th and 25th percentiles of the income distribution from the ACS (providing national representativeness) and take the difference between the two. We estimate the relationship in logs rather than in U.S. dollars because it captures how price differences flatten out at higher incomes.[6] For 2004 this yields a gap of about 5.1 percentage points.[7] We repeat this quarter by quarter to trace the 75–25 price gap over time in figure 2.

2. The 75–25 price gap for consumer packaged goods, 2004–22

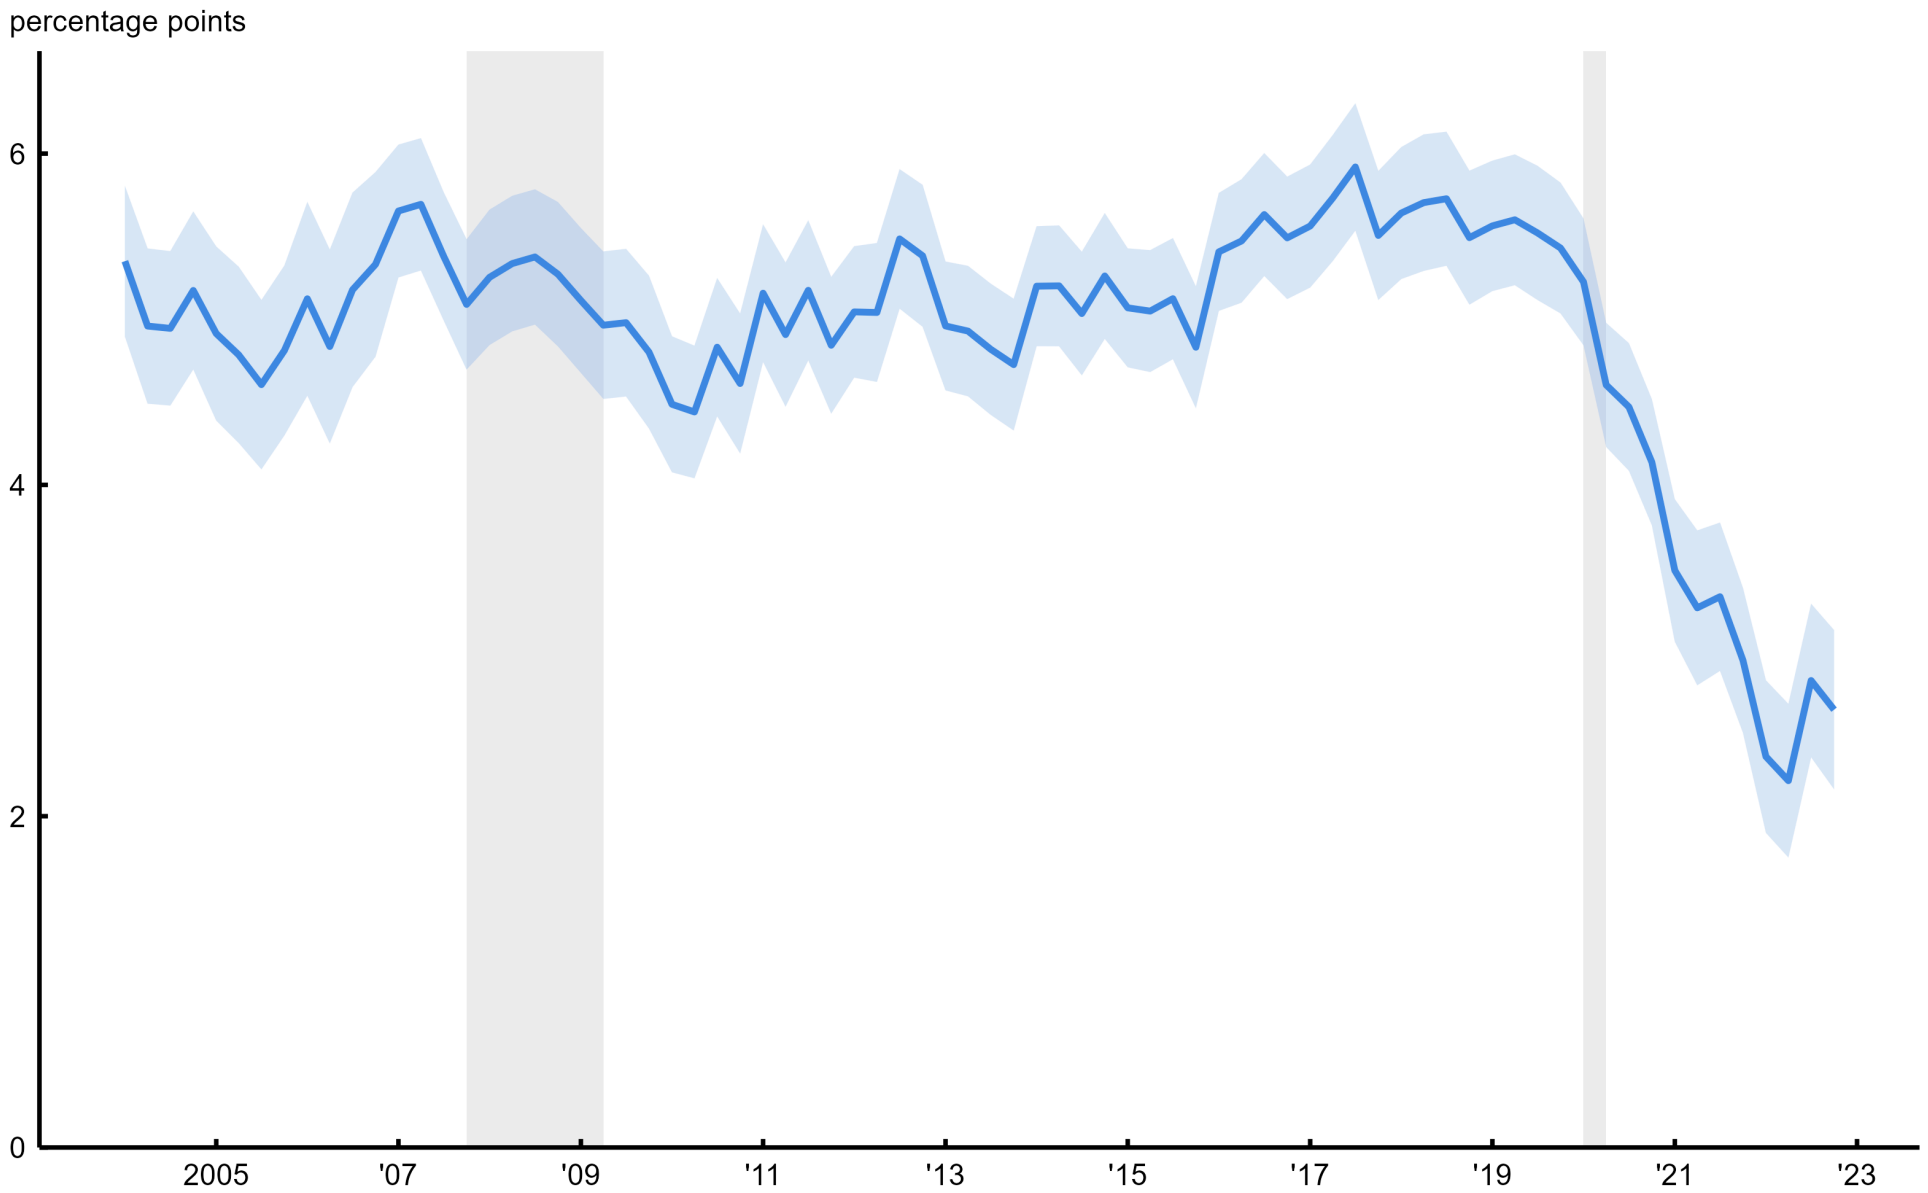

Notes: The figure shows the estimated price difference for consumer packaged goods between the households in the 75th and 25th percentiles of the income distribution, expressed in percentage points, for each quarter in the sample (dark blue line). The estimate is obtained by regressing the percentage deviation from the national category average against the logarithm of the income group midpoint, separately by quarter, and then predicting the difference in price deviations between the 75th and 25th percentile household incomes from the *American Community Survey*. The light blue shaded band (around the dark blue line) is the 95% confidence interval based on standard errors clustered by household. See the text for further details. The gray vertical bands mark recessions as determined by the National Bureau of Economic Research.

Sources: Authors' calculations based on data from the NielsenIQ Consumer Panel Service and U.S. Census Bureau, *American Community Survey*.

Our 75–25 price gap is a useful measure but has some limitations. On the positive side, because prices are measured as deviations from category averages, they are unaffected by changes in the aggregate level of inflation and we can focus on the relative price differences over time. Moreover, scaling by the ACS income gap is better than fixing U.S. dollar thresholds for the comparison groups (as NielsenIQ intervals for the income bins are for the most part constant through time). A household earning $20,000 in 2005 is a very different household from one earning $20,000 in 2022, since both inflation and real income growth have pushed the typical household's earnings up; thus, any fixed cutoff drifts toward the bottom of the distribution over time. Anchoring to the same percentiles each year keeps the comparison at a fixed relative position in the distribution.[8]

Figure 2 plots the 75–25 price gap over time. The gap increased from 5.35 percentage points in 2004:Q1 to 5.65 percentage points in 2007:Q1 and then decreased slowly to 4.49 percentage points by 2010:Q1. It held mostly flat through the 2010s before falling sharply after 2020:Q1, down to 2.36 percentage points by 2022:Q1. The recent drop and the 2007–10 drop both coincide with economic downturns, when households shop more carefully and trade down toward cheaper options and stores (Aguiar et al., 2013; Nevo and Wong, 2019; and Coibion et al., 2015). Overall, the 75–25 price gap shows little correlation with gross domestic product (GDP) growth (the correlation is −0.19) or with unemployment (essentially zero). It is negatively correlated with inflation (ρ=−0.64), though this relationship is driven mostly by the post-2020:Q1 period, when inflation rose and the 75–25 price gap fell at the same time. The correlation between the 75–25 gap and inflation is 0.04 if we restrict our analysis to the 2004–19 period.

What explains the post-2020:Q1 decline in the 75–25 price gap?

The decline in the 75–25 price gap for consumer packaged goods after 2020:Q1 was particularly steep, and given its recency and size it is natural to ask what drove this decrease and what it means. We proceed in two steps. First, we ask whether the 75–25 price gap fell because prices converged *within* categories or because households reallocated spending *across* categories. The answer is almost entirely within-category: For the same categories of goods, the households at the 75th and 25th percentiles of the income distribution now pay more similar prices. Second, we ask what that convergence implies for inflation. We show that the gap can narrow only if the households with lower incomes face higher inflation than the ones with higher incomes—a result that holds in theory and that our data confirm.

Did the 75–25 price gap fall within categories or across them?

The 75–25 price gap is a weighted average of the price gaps between the two income groups in each product category, where the weight is how much people spend on that category. This structure means the headline number can move for two very different reasons, and the distinction matters for what the decline tells us. Before fully exploring the two different reasons, we provide a simple base example: Imagine households only purchase two types of goods—coffee (a category with a large price gap between the two income groups) and toilet paper (a category with a small price gap between them). In this example, the overall 75–25 price gap is the weighted average of these two gaps, with weights equal to their expenditure shares.

The first reason the 75–25 price gap could change is within-category convergence. Take the category of coffee: If the prices the households with higher and lower incomes pay for coffee move closer together, then coffee's own price gap narrows and the overall number falls. For the same kind of good, the two groups now pay more similar prices. The second reason the gap could change is across-category convergence and has nothing to do with prices of the same type of good moving closer to each other. Categories differ in how large their price gap between the income groups is. So, if households simply shift spending toward the low-gap categories (e.g., shifting from coffee to toilet paper), the average falls even though no single category's price gap has moved. The makeup of the basket changed, not the prices of the items within it.

These two channels carry different implications. Convergence of prices within categories indicates the price difference between households with higher and lower incomes is decreasing. The reallocation of spending (by either income group) across categories means only that spending patterns shifted—the underlying price gaps for particular types of goods between income groups could be the same as before. Distinguishing between these two channels tells us where to look further for an explanation of why the overall 75–25 price gap is converging (or diverging).

To understand if the decrease in the 75–25 price gap is within- or across-category, let G_t^c denote the price gap between households at the 75th and 25th percentiles of the income distribution for category c in quarter t and let s_t^c denote the expenditure share of that category in that quarter. The overall 75–25 price gap, G_t, is their expenditure-weighted average,

$$G_t = \sum_c s_t^c G_t^c.$$

The gap's change between any arbitrary periods $t-k$ and t then splits into two terms:

$$G_t - G_{t-k} = \underbrace{\textstyle\sum_c s_{t-k}^c \left(G_t^c - G_{t-k}^c \right)}_{\text{(a) within-category}} + \underbrace{\textstyle\sum_c \left(s_t^c - s_{t-k}^c \right) G_t^c}_{\text{(b) across-category}},$$

where the first term (a) holds expenditure shares fixed and asks how much the gap moved within categories and where the second term (b) holds the within-category gaps fixed and asks how much households shifted spending across categories. Using this equation, we can decompose the change over the period 2020:Q1–2022:Q4, when the gap fell by 2.6 percentage points, into the two components. According to our analysis, the decline in the 75–25 price gap during this span (when the Covid pandemic started) is entirely due to within-category changes; the within-category term accounts for slightly more than the entire decline (about 102%), while spending reallocation is small and slightly offsetting (about –2%).

The convergence in the prices paid by households at the 75th and 25th percentiles of the income distribution was broad-based, not concentrated in a handful of goods. Figure 3 plots, for the 20 largest categories (as defined by their 2020–22 expenditure shares) how each category's contribution to the 75–25 price gap moved from 2020:Q1 (hollow dot) to 2022:Q4 (filled dot).[9] For the vast majority of categories, the filled dot sits to the left of the hollow dot, indicating its contribution fell. Everyday staples—e.g., fresh fruit, fresh bread, and milk—are among the largest contributors to the decline. To further decompose the movement between the two periods for every category, we split each contributor into two pieces. The first one is a within-category piece (blue arrow), which tracks the change in the category's own price gap; it is negative when the gap falls, or $G_t^c < G_{t-k}^c$. The second is an across-category piece (red arrow), which tracks the change in the category's

expenditure share. Since gaps for most categories are positive ($G_t^c > 0$), the across-category piece is negative when the expenditure share falls in the category ($s_t^c < s_{t-k}^c$). The within-category arrow is responsible for most of the decline: The across-category arrows tend to be short, while the within-category arrows tend to be longer and point to the left.

3. Drivers of the 2020–22 decline in the 75–25 price gap for consumer packaged goods, by product category

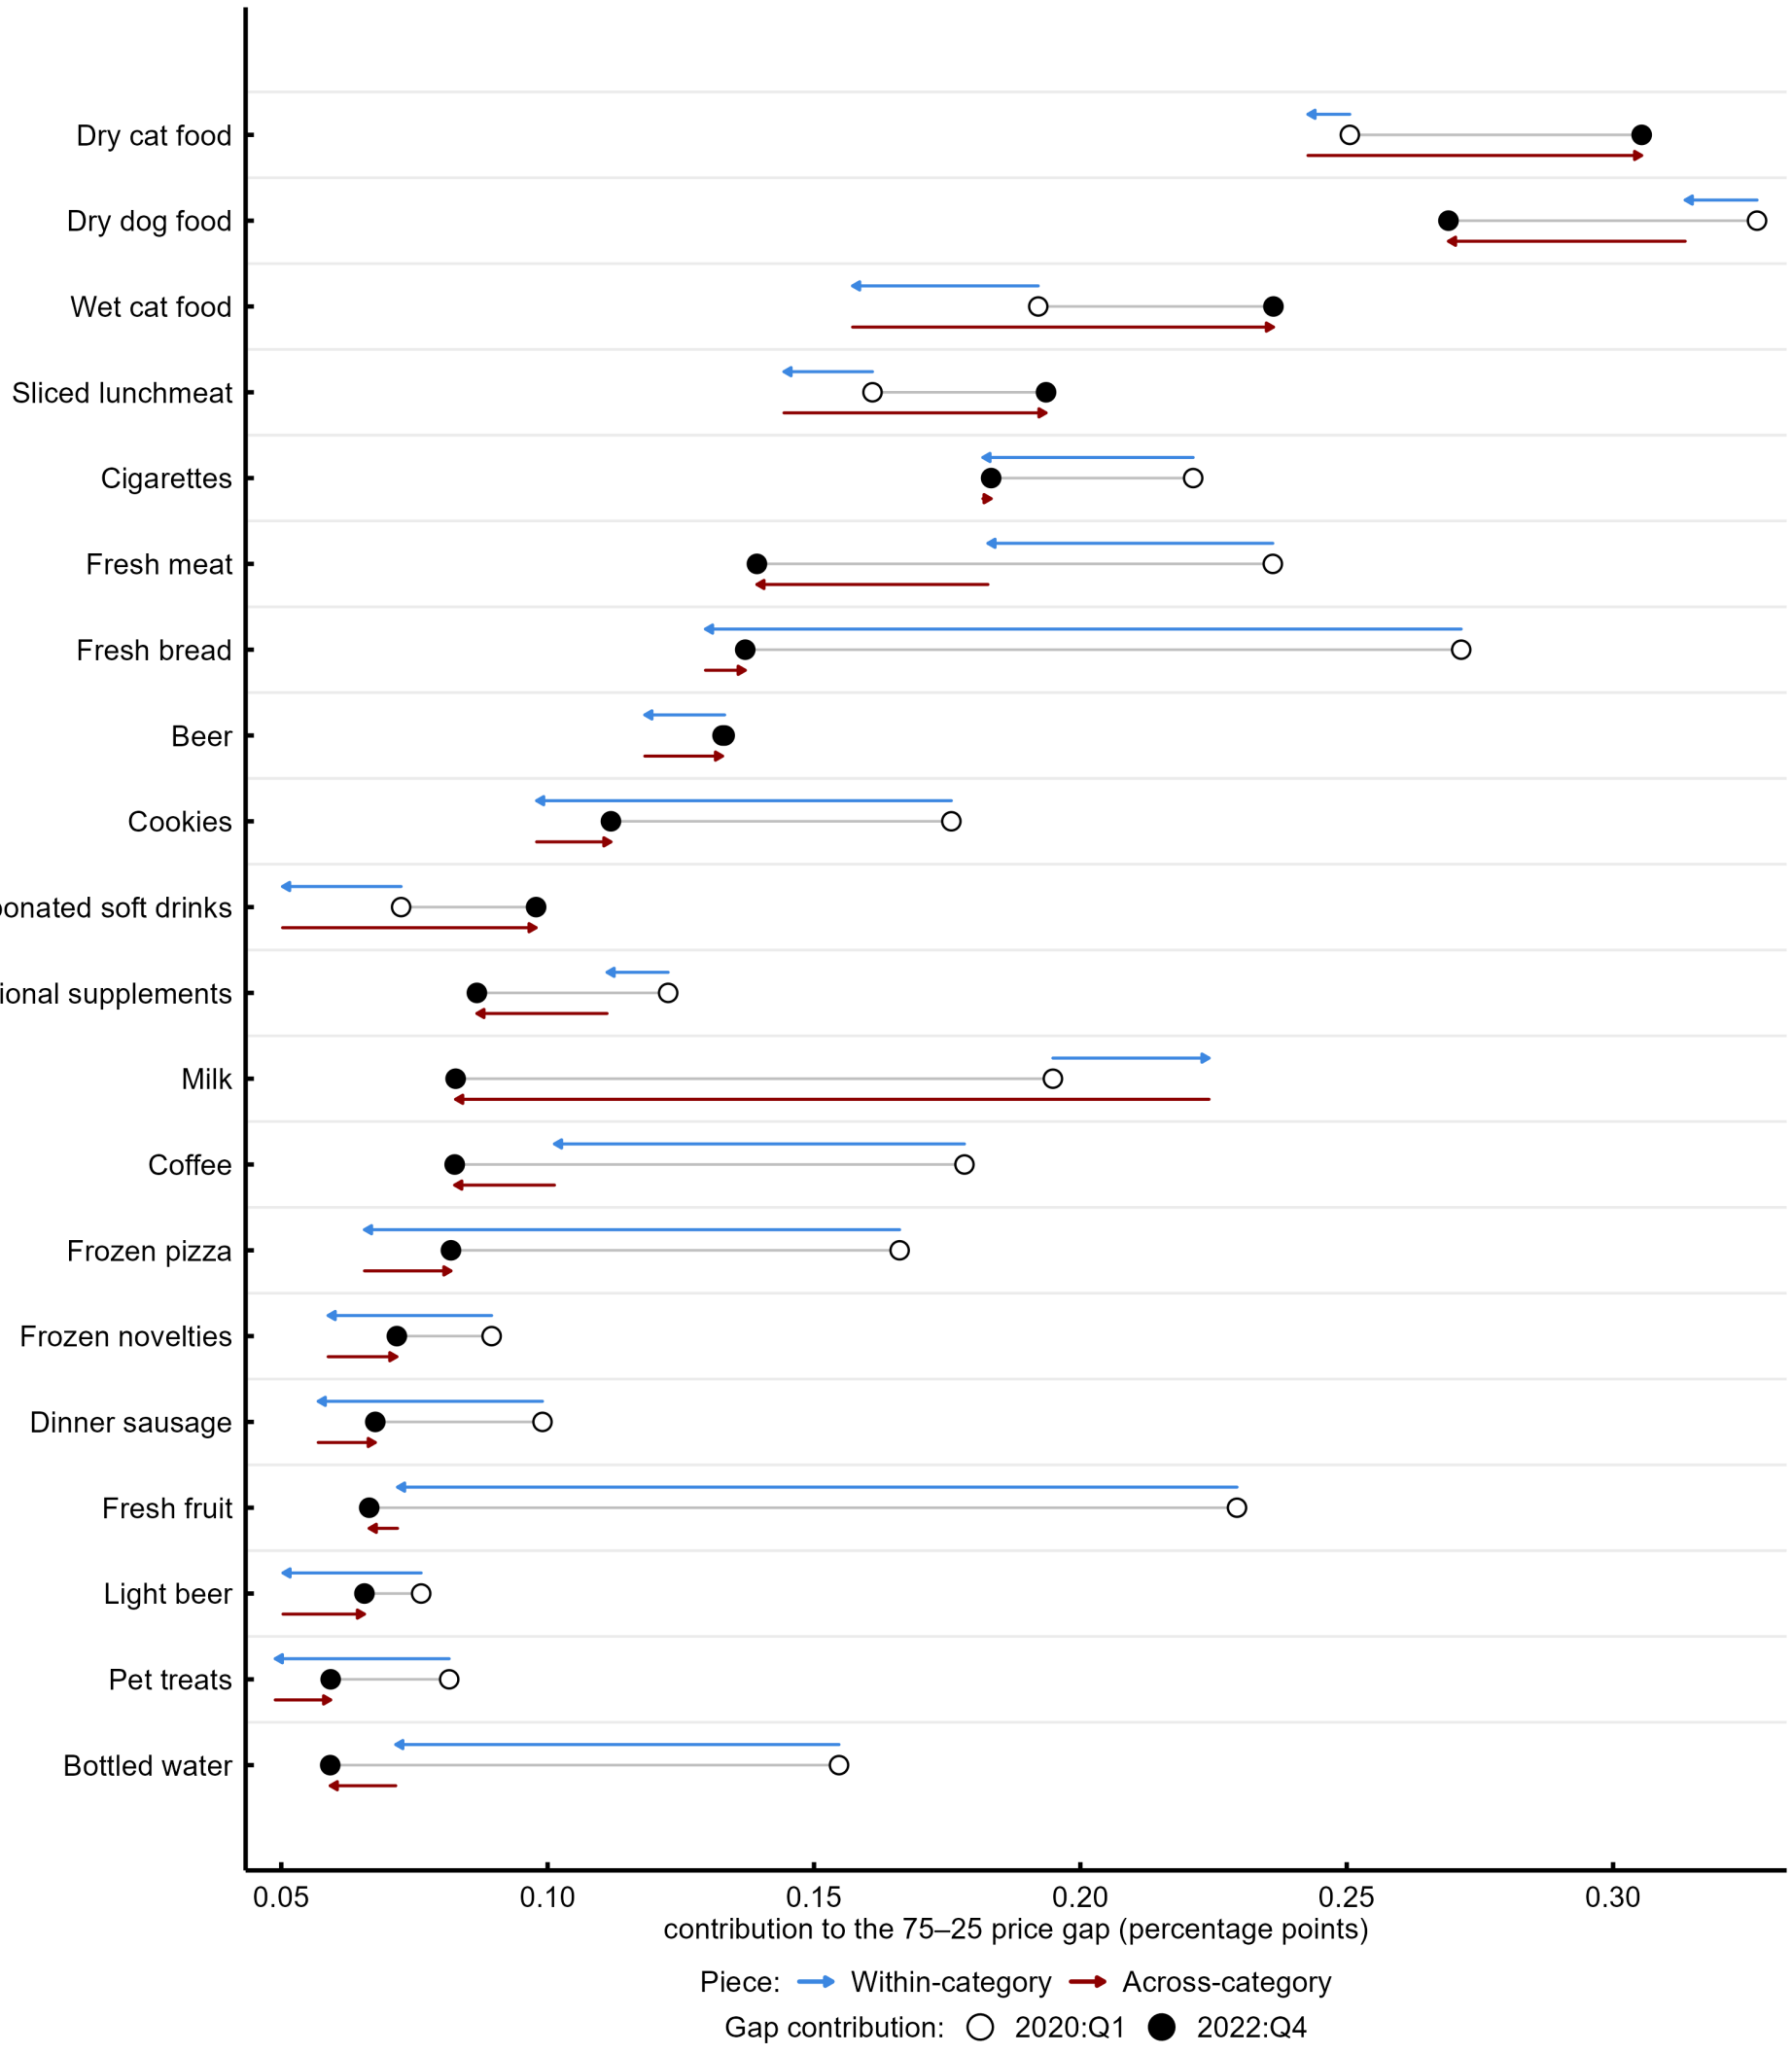

Notes: The figure shows, for the 20 largest categories of consumer packaged goods by their 2020–22 expenditure shares, each category's contribution to the aggregate 75–25 price gap—i.e., the price gap between households at the 75th and 25th percentiles of the income distribution. For each category the hollow dot is its contribution at the start (2020:Q1) and the filled dot its contribution at the end (2022:Q4); a category's contribution is its expenditure share times its own 75–25 gap, so the dots sum across categories to the aggregate gap. The move from start to end splits into two pieces, shown as arrows: The within-category piece (blue) holds the category's base-period expenditure share fixed and reflects the change in its own gap, and the across-category piece (red) holds its end-period gap fixed and reflects the change in its expenditure share. The two pieces sum exactly to the total move (hollow to filled dot). The categories are ordered by their end-period contribution, with the category with the largest contribution at the top. All quantities are in percentage points and weighted by the NielsenIQ household projection factor times expenditure.
Sources: Authors' calculations based on data from the NielsenIQ Consumer Panel Service and U.S. Census Bureau, *American Community Survey*.

Is inflation related to the narrower gap?

In the previous subsection we document that the 75–25 price gap fell sharply after 2020:Q1 because of the within-category convergence in prices. How is this related to inflation? A large literature, surveyed in Jaravel (2021), documents that inflation differs systematically across the income distribution (Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019; and Argente and Lee, 2021); the goal of this analysis is to connect that literature to our price gap measure.

We now decompose the price gap for a given category (G_t^c) in a slightly different way than before to link it to inflation. We denote by p_{ct}^R and p_{ct}^P the expenditure-weighted average price faced by households with higher and lower incomes (which we continue to represent using the 75th and 25th percentiles of the income distribution, respectively) and by $\bar{p}_{ct}$ the expenditure-weighted national average price in that category. The gap is the difference between the two deviations:

$$G_t^c = \underbrace{\frac{p_{ct}^R - \bar{p}_{ct}}{\bar{p}_{ct}}}_{\substack{\text{75th percentile} \\ \text{household deviation}}} - \underbrace{\frac{p_{ct}^P - \bar{p}_{ct}}{\bar{p}_{ct}}}_{\substack{\text{25th percentile} \\ \text{household deviation}}} = \frac{p_{ct}^R - p_{ct}^P}{\bar{p}_{ct}}.$$

The identity applies both category by category and when aggregating across categories.[10] The inflation logic that follows therefore holds category by category as well as in aggregate.

We relegate the slightly cumbersome algebra to the appendix, but show that for the price gap for a given category to fall between 2020:Q1 and 2022:Q4 (or more generally between $t-k$ and t), households with lower incomes must have faced higher inflation than households with higher incomes over that window:

$$G_t^c - G_{t-k}^c < 0 \Leftrightarrow \pi_c^R < \pi_c^P,$$

where π_c^R and π_c^P represent the percentage change in prices (i.e., what we call, somewhat loosely, inflation) for a given category for households with higher and lower incomes, respectively.[11] The intuition is simple: The price gap is the percentage point difference between the two price levels. If inflation is the same across the two income groups, the two price levels change in equal proportions and the gap remains constant. Since households with lower incomes pay lower prices, the change in the gap can only be negative if their prices rise faster than those for households with higher incomes. In other words, a narrowing gap means households with lower incomes are experiencing faster price increases.

We can make these calculations more concrete using NielsenIQ data. We price a fixed basket of goods for two representative households—one at the 25th percentile of the income distribution and another at the 75th percentile, holding each group's basket at its 2019 quantities and asking what each basket of goods cost in 2020:Q1 and in 2022:Q4. To create the basket of goods, we use 2020:Q1 data and interpolate the expenditures for our two representative households.[12] We allow both income groups to have different consumption baskets (e.g., households with higher incomes may spend more of their money on coffee than those with lower incomes by buying more of the same coffee or buying different brands of coffee). We then calculate how much households would spend in 2022:Q4 by using that period's prices. For simplicity's sake, we only allow prices to differ between the two dates.

Again, take coffee, for instance. The price of a 12-ounce bag rose from about $4.11 in 2020:Q1 to $6.26 in 2022:Q4 for the lower-income household (a 52% increase), but from $4.56 to $6.65 for the higher-income household (a somewhat smaller 46% increase). The lower-income household paid less for coffee to begin with, but its price climbed faster in percentage terms, so the prices for coffee paid by the two representative households ended up closer. Fresh bread shows the same pattern: A 20-ounce loaf went from $2.80 to $3.88 for the lower-income household (up 39%) versus $3.06 to $4.04 for the higher-income one (up just 32%). Just as in figure 3, the pattern holds across everyday staples. Even where the two rates were close, the lower-income household's was rarely lower—e.g., a two-liter bottle of soda rose 39.3% for the lower-income household versus 38.5% for the higher-income one, and similarly, a 12-pack of beer was up 15.8% for the lower-income household versus 14.5% for the higher-income one.[13]

After being summed across the whole basket of goods, these category changes give us the difference in the total basket cost between 2020:Q1 and 2022:Q4 in figure 4—namely, what each group's 2019 quantities of the 50 categories of consumer packaged goods cost at the prices it faced in 2020:Q1 versus 2022:Q4. The lower-income basket cost rose 39.7%—from about $1,459 in 2020:Q1 to $2,038 in 2022:Q4. And the higher-income basket rose 36.5%—from about $1,548 in 2020:Q1 to $2,113 in 2022:Q4. To sum up, the lower-income household (at the 25th percentile of the income distribution) experienced inflation in consumer packaged goods that was about 3.2 percentage points higher than the higher-income household (at the 75th percentile), narrowing the 75–25 price gap over the 2020–22 period.

4. Cost of a fixed 2019 basket of consumer packaged goods in 2020:Q1 versus 2022:Q4: Households in the 25th versus 75th percentile of the income distribution

	2020:Q1 ($) (1)	2022:Q4 ($) (2)	Change ($) (3)	Change (%) (4)
25th percentile	1,459	2,038	579	39.70
75th percentile	1,548	2,113	565	36.50

Notes: Each row prices a single representative household's fixed basket—its 2019 average annual quantities across the 50 categories of consumer packaged goods, in each category's unit of measure—at the fitted per-unit prices it faced in 2020:Q1 and in 2022:Q4. Quantities are held fixed across the two dates, so columns 3 and 4 reflect price changes alone. The households at the 25th and 75th percentiles are recovered by interpolating each category's quantities and prices across income bins (or groups) to the 25th and 75th *American Community Survey* income percentiles (held fixed at their 2020 values). The quantities are averaged over all households in the income group, including those who bought none of a packaged good category. All amounts are weighted by the NielsenIQ household projection factor.
Sources: Authors' calculations based on data from the NielsenIQ Consumer Panel Service and U.S. Census Bureau, *American Community Survey*.

What is even more interesting is that the decline in the 75–25 price gap represents a convergence in *levels*, but driven by the lower-income households' prices rising faster in *percentage* terms. In percentage terms, the lower-income household faced faster inflation, 39.7% versus 36.5% for the higher-income household, as we just discussed. In nominal U.S. dollars, though, the two increases are nearly identical, about $579 for the low-income household and $565 for the higher-income one. How can the percentage increases differ that much while the dollar increases barely do? The lower-income household started from a cheaper basket—about 5.7% cheaper in 2020:Q1. A larger percentage increase on a smaller base lands at almost the same U.S. dollar value.[14] This is consistent with Sangani (2026), who documents that cost shocks pass through to retail prices roughly one for one in U.S. dollar terms rather than in percentage terms, so that cheaper varieties see larger percentage price increases.[15]

In figure A3 in the appendix, we break the cost of the fixed 2019 basket down by product category to understand if the increase in figure 4 comes from a narrow set of goods or if it is broad-based. All categories but one (paper towels) show cost increases. What is most interesting about figure A3 is that the dollar-versus-percentages story from figure 4 holds in most categories. The dollar increase is nearly identical for lower- and higher-income households—e.g., orange juice costs $3.97 more in 2022:Q4 than in 2020:Q1 for both groups, while a frozen poultry entree costs $12.15 and $11.25 more for lower- and higher-income households, respectively. Because lower-income households typically start from a lower base, these similar U.S. dollar increases translate into larger percentage increases. This points to cost pass-through as a likely driver of the convergence in prices. For instance, retailers raised prices by a similar dollar amount per unit regardless of who buys the type of packaged good.

Conclusion

In U.S. markets for consumer packaged goods, the households with higher incomes (i.e., those at the 75th percentile of the income distribution) pay more than the households with lower incomes (i.e., those at the 25th percentile) for the goods in the same category, but the former income group pays less of a premium than before. The 75–25 price gap has shrunk by more than half since 2004, and most of that decline came after 2020:Q1. It is a within-category story: For the same kinds of goods, households with higher and lower incomes now pay more similar prices. The 75–25 price gap did not shrink because those with greater means rearranged their baskets toward cheaper categories. This gap can then only close when lower-income households' prices rise faster—and in the post-2020:Q1 window they did, by about 3 percentage points. Perhaps paradoxically, a decrease in pricing inequality, which at face value some may consider desirable, means higher inflation for lower-income households than for higher-income ones. While we treat these patterns as descriptive and leave their mechanisms and welfare consequences for future work, our findings are consistent with cost shocks being passed through to consumers one for one in U.S. dollar terms.

All errors in this article are our own. Our own analyses and calculations are based in part on NielsenIQ Consumer Panel Service data for all categories for the 19-year period ending in 2022 for the total U.S. The conclusions drawn from the NielsenIQ data are ours and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein or in developing, reviewing, or confirming the research approaches used in connection with this article.

Notes

[1] This statement may run against intuition, since in many other settings households with lower incomes pay more. For example, Attanasio and Pastorino (2020) document that quantity discounts are sizable for basic staples, favoring bulk buyers, and Orhun and Palazzolo (2019) show that liquidity-constrained, low-income households buy on quantity discount less often than higher-income ones, forgoing savings that the latter capture. Lower-income households may also face higher prices because they live farther away from discount retailers. Allcott et al. (2019) document that low-income neighborhoods are more likely to be "food deserts"—areas with fewer supermarkets and, consequently, higher prices—though they show that this supply-side variation explains only a small share of consumption differences across income groups. With that said, lower-income households can and do adopt shopping behaviors that often result in their paying less than their higher-income counterparts for the same categories of goods: More intensive shopping and search can lower the price paid for identical goods (Aguiar and Hurst, 2007; and Kaplan and Menzio, 2015), and households can save through outlet choice, bulk buying, sales, and generic brand item purchases (Griffith et al., 2009). Like any of these shopping behaviors and options, bulk buying can be chosen by either higher-income or lower-income households so that they end up paying lower prices relative to what others in different segments of the income distribution pay for the same categories of goods (but the specific categories with better prices paid may differ by income group).

[2] We exclude transactions with missing prices, and we drop the top and bottom 1% of prices for each product category-year to mitigate the influence of outliers in our results. We also drop households whose average monthly expenditure falls in the top or bottom 2.5% of the household income distribution, since these extreme values most likely reflect outliers or misreporting—i.e., institutional or business buyers or data-entry errors—rather than genuine consumption.

[3] It is worth noting that our product categories leave out random-weight perishables sold by weight at the counter (loose produce, deli, and butcher-counter meat), which NielsenIQ slowly phased into the sample starting in 2007, making them hard to track over time.

[4] Our figure 1 resembles figure 2 of Broda et al. (2009), with one key difference: They focus on within-product differences, while we look at the average price paid across all products in a category.

[5] A regression is a statistical process that measures the degree of correlation between two variables—an independent (predictor) variable and a dependent (response) variable—while holding constant the other independent variables; the estimated coefficient from a regression represents the mean change in the dependent variable for a one-unit change in the independent variable. For the income distribution's topmost bin, which is open-ended, we assign the average income of households in that bin from the *American Community Survey*. We exclude the lowest-income bin (households earning less than or equal to $4,999 per year), since we focus on the price gap in the middle of the distribution and the nonmonotonicity of these households with low incomes in the sample flattens the slope of the price difference.

[6] When we tried a level relationship instead, it understated the price gap between the households in the 75th and 25th percentiles of the income distribution in most years.

[7] This is slightly lower than the 5.9 percentage points we read directly off figure 1 because the 75th percentile sits at the lower end of its income bin rather than at the midpoint. In other words, reading the price gap from the income bins somewhat overstates this difference.

[8] There are two limitations of this approach: It relies on income bin midpoints to estimate price differences, which introduces measurement error in the regression's independent variable, making the estimates less precise; and it imposes a parametric (log-linear) form on the relationship between income and price differences, which risks misspecifying the regression model if the true relationship is not log-linear. We believe, though, that a log-linear relationship is appropriate given our results in figure 1.

[9] In our notation, the hollow dot plots the expenditure share multiplied by the gap for a given category in 2020:Q1 (or $s^c_{t-k} G^c_{t-k}$) while the filled dot plots the same, but in 2022:Q4 (or $s^c_t G^c_t$).

[10] Read p^R_t, p^P_t, and $\bar{p}_t$ as overall expenditure-weighted price levels and G_t as the aggregate price gap.

[11] We use the term inflation loosely here because our measure is the percentage change in prices for a given category, aggregated across categories using initial expenditure weights. It differs from headline measures of inflation such as the Consumer Price Index or the Personal Consumption Expenditures Price Index; those inflation indexes cover other goods and services outside of consumer packaged goods and aggregate categories using different weighting schemes.

[12] The NielsenIQ panel records income in bins, not a household exactly at a given percentile. As we do throughout this article, we recover the percentile household by regressing the quantity (or price) of each category on log income across bins and evaluating the fit at the 25th and 75th ACS income percentiles; both percentiles lie inside the range of the NielsenIQ

data, so this interpolates rather than extrapolates. Quantities are the 2019 average annual amount purchased in each category's unit of measure, averaged over all households in the income group, including nonbuyers; prices are the fitted price per unit in each quarter. The basket spans the 50 categories used throughout this article.

[13] Authors' calculations based on data from the NielsenIQ Consumer Panel Service and U.S. Census Bureau, *American Community Survey*.

[14] As a check on the method, this fixed-basket approach reproduces the headline 75–25 price gap by an entirely separate route. Pricing a single common basket at each group's fitted prices, the higher-income household's basket costs about 5.5% more than the lower-income household's basket in 2020:Q1—essentially the 75–25 price gap of 5.2 percentage points that the regression underlying figure 2 yields for the same quarter.

[15] It is also consistent with Sullivan (1985), who documents that for the cigarette industry, shocks affecting the cost per unit (in this case, excise taxes) pass through in U.S. dollar terms (increasing the price to $1.09 from $1.00) rather than proportionally, likely generating larger price increases among cheaper products.

Opinions expressed in this article are those of the author(s) and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.