Inflation Resilience*

Jennie Bai† Francesco D'Acunto‡ Xiao Yin§

July 2025

Preliminary and Incomplete. Comments Welcome

Abstract

Unique US-wide data from over 200M transactions by 11M consumers across 5,000 brands from 2020 to 2022 reveal wide spatial and industry variation in the extent to which consumers maintain consumption quality amid rising inflation (*inflation resilience*). A 1 ppt rise in local inflation lowers the local share of high-quality spending by 0.68 ppt on average but with standard deviations of 1.6 ppt across counties and 1.4 ppt across industries. We first confirm that higher-income counties display lower resilience because of a higher pre-inflation high-quality share. We then lever our unique mapping of brands, transactions, and retailers to show that the local retail market structure matters: both lower local brand variety and higher retail concentration, ceteris paribus, lead to lower resilience in the face of the same increase in local inflation. Inflation resilience helps explain behavior beyond consumption: it is not higher local inflation per se but lower inflation resilience keeping inflation levels constant that explains the 2024 US Presidential election shift from Biden to Trump, suggesting agents despise the decline in consumption quality that often accompanies inflation surges rather than inflation itself.

JEL Classification: D12, D72, E21, E31, L11, R12
Keywords: micro-to-macro, consumption, monetary economics, IO, trade, economic geography, economic fluctuations, political economy, big data, LLMs.

*We thank Andrea Savio for excellent research assistance. All errors are our own.
†McDonough School of Business, Georgetown University, and the National Bureau of Economic Research, Email: Jennie.Bai@georgetown.edu.
‡McDonough School of Business, Georgetown University, Email: francesco.dacunto@georgetown.edu.
§Department of Economics and School of Management, University College London, Email: xiao.yin@ucl.ac.uk.

Inflation Resilience

Abstract

Unique US-wide data from over 200M transactions by 11M consumers across 5,000 brands from 2020 to 2022 reveal wide spatial and industry variation in the extent to which consumers maintain consumption quality amid rising inflation (*inflation resilience*). A 1 ppt rise in local inflation lowers the local share of high-quality spending by 0.68 ppt on average but with standard deviations of 1.6 ppt across counties and 1.4 ppt across industries. We first confirm that higher-income counties display lower resilience because of a higher pre-inflation high-quality share. We then lever our unique mapping of brands, transactions, and retailers to show that the local retail market structure matters: both lower local brand variety and higher retail concentration, ceteris paribus, lead to lower resilience in the face of the same increase in local inflation. Inflation resilience helps explain behavior beyond consumption: it is not higher local inflation per se but lower inflation resilience keeping inflation levels constant that explains the 2024 US Presidential election shift from Biden to Trump, suggesting agents despise the decline in consumption quality that often accompanies inflation surges rather than inflation itself.

1 Introduction

Inflation erodes households' real purchasing power, forcing consumers to adjust either the quantity or the quality of goods and services they consume. In periods of high inflation, households may respond by buying fewer items of a given quality or by switching to lower-quality but lower-price alternatives to keep their consumption quantity unchanged. This behavior, often termed "trading down," has been documented empirically (Jaimovich et al., 2019; Argente and Lee, 2021; Jaravel, 2021).

In this paper, we define *inflation resilience* as the extent to which consumption quality stays stable in the face of rising inflation in a location or industry, that is, the extent to which consumers do not trade down. Using novel transaction-level data covering over 200 million purchases made by 11 million consumers spanning 5,000 distinct brands across all 50 U.S. states from 2020 until 2022, we document, to the best of our knowledge for the first time, substantial variation in inflation resilience across space and industries and we investigate its drivers and consequences.

Because our data span a wide range of product categories and brands, we can construct homogeneous measures of consumption quality across space and industries. Our baseline measure of quality classifies transactions into high-, medium, and low-quality tiers based on brands' average price per transaction at the county or industry level. We then compute the share of consumption expenditure and transactions in each quality tier for all county-months and industry-months in our sample. For robustness purposes, we also use non-price-based criteria to capture brands' quality, such as perceived brand luxury rankings assigned by ChatGPT. We relate these consumption quality shares to changes in local monthly inflation to estimate inflation resilience at the county and industry levels.

In our baseline county-industry-level panel regression specification, which includes a comprehensive set of fixed effects and demographic controls, a one percentage point (ppt) increase in local inflation is associated with a 0.68 ppt decrease in the share of expenditures on high-quality brands relative to low-quality brands, which confirms that, on average, US households

trade down in the face of rising inflation. Crucially, we show that this effect reflects active substitution by households rather than a purely mechanical price index effect,[1] because higher inflation leads consumers to increase the number of purchase transactions of lower-quality goods relative to higher-quality goods keeping constant consumption quantity.

Having confirmed that inflation resilience is on average negative, we move on to document its stark heterogeneity across US counties and industries. During the 2021-2022 inflation surge, some US counties saw almost no change in consumption quality whereas others experienced large declines, even when facing similar changes in local inflation. The standard deviation of inflation resilience across counties is about 1.6 ppt with an interquartile range (IQR) of 1.7 ppt: less resilient counties had consumption quality declines roughly three times larger than more resilient counties that faced the same change in local inflation. We also detect substantial heterogeneity across industries: for 3-digit NAICS sectors, the standard deviation is 1.4 ppt and IQR approximately 0.8 ppt.

This susbtantial heteorogeneity of inflation resilience across space and industries motivates us to dig deeper into understanding what determines a location's or industry's inflation resilience. We first consider demographic characteristics. Theory and existing evidence on trading down suggest that, because high-income households generally consume higher-quality goods, they have more scope to reduce consumption quality keeping their quantity unchanged; by contrast, lower-income households, who already spend mostly on lower-quality goods, have less scope to trade down (Argente and Lee, 2021). Our evidence supports this mechanism across space: counties with higher median income tend to have lower inflation resilience. For example, counties in the top third of the income distribution saw a quality share decline about 0.22 ppt larger than counties in the bottom third that face the same increase in local inflation (a nearly 50% larger effect). Crucially, we show directly that this heterogeneity by local income levels is fully explained by the pre-inflation-surge level of the local low-quality consumption share: once we control for this share in our regressions, the conditional

[1] If the prices of lower-quality goods rise relatively more than those of higher-quality goods, a nominal expenditure-based quality share could fall even without consumers actively changing their bundle.

association between income and inflation resilience becomes economically and statistically insignificant. We further corroborate this mechanism by showing that all the other demographic characteristics we observe, such as local education levels, age distributions, and racial composition, among others, do not display any significant correlation with inflation resilience once we control for income and the pre-inflation-surge low-quality consumption share.

We then exploit the mapping between transactions, retailers, brands, and locations in our data to assess the potential role for the local retail market structure in explaining inflation resilience. First, we conjecture that when consumers face more local brand variety, and hence many alternatives within each quality tier, they could adjust at the intensive margin by switching to same-quality brands that might be relatively cheaper or slower to increase their retail prices, rather than at the extensive margin by downgrading to lower-quality brands. This conjecture is consistent with recent evidence such as Cavallo and Kryvtsov (2024), who document that consumer exploitation of within-store discounts partly helps them reduce their bundles' inflation, whereas the tendency to switch to cheaper products ("cheapflation") exacerbates it. We find robust evidence for this mechanism: counties with a greater number of distinct brands available before the inflation surge display higher inflation resilience, ceteris paribus. Specifically, counties in the top quartile of brand variety display about 37% higher inflation resilience than those in the bottom quartile. Within industries, higher brand variety leads to up to 75% higher inflation resilience.

Second, our data allow us to assess the potential role of local retail concentration in explaining inflation resilience. Intuitively, in locations with fragmented retail markets where multiple retailers sell goods from the same brands, consumers have the possibility to shift across retailers that use different pricing strategies in times of inflation to keep consuming in the same quality tier by chasing lower price increases. By contrast, in locations with high sales concentration where a few large retailers dominate transactions and impose uniform pricing strategies (e.g., see DellaVigna and Gentzkow (2019)), consumers should find it harder to maintain the quality of the brands they consume without facing higher prices. Consistently,

we find that inflation resilience is economically and statistically lower in counties in which retail transactions are more concentrated relative to those in which transactions are more dispersed across retailers.

An important concern with our baseline analysis is endogeneity. Reverse causality is a concern because consumers' trading down behavior increases the local demand of low-quality goods, whose larger relative price increases could lead to a higher reading of local overall inflation. Moreover, omitted variables that move hand in hand with local inflation might be what actually drives consumers' assessment of the trade-off between consumption quality and quantity in times of growing inflation. We tackle these concerns in two ways. First, we note that our baseline estimates of inflation resilience are systematically larger in our most restrictive specifications, which control for county-by-month effects, county-by-industry fixed effects, and a host of observable local characteristics, and hence absorb directly many potential confounding factors. This evidence suggests that potential omitted demand or supply factors, if anything, were biasing our baseline OLS estimates toward zero.

Second, we implement an instrumental-variable (IV) strategy to aim to isolate quasi-exogenous variation in local inflation in a shift-share instrument setting. Specifically, we exploit two ingredients: an external cost shock due to price increases in manufacturing good inputs imported from China over the period 2021−2022 ("shift") and the pre-2021 exposure of each US state's economy to manufacturing ("share"). The idea is that input price increases imposed outside the US, which are largely driven by supply conditions in China, led to higher production costs and inflationary pressure in US locations that had a larger manufacturing sector and hence were more exposed to such external shocks, irrespective of local demand conditions. Using the interaction between these two variables as our instrument for local inflation, we find that this instrument is relevant, with first-stage F-statistics of approximately 80 across alternative specifications. This instrument would satisfy the exclusion restriction as long as Chinese input cost shocks did not affect local households' consumption quality choices through direct channels other than local inflation. Using a two-stage least squares estimator,

we find that higher instrumented local inflation leads to a more pronounced trading-down effect on average, corroborating our OLS regression results.

In the last part of the paper, we aim to shed light on whether inflation resilience is perceived as a negative or positive outcome by consumers. This question is inspired by our evidence, which suggests a trade-off. On the one hand, and consistent with the standard interpretation of trading-down behavior in macroeconomics, the possibility of trading down is an advantage for consumers, who can shield their bundles and consumption quantities from the negative effects of inflation by adjusting along the quality margin. Under this interpretation, for instance, the fact that higher-income consumers can trade down while low-income consumers cannot is interpreted as an advantage for high-income consumers in times of rising inflation. At the same time, though, a large literature in social psychology and marketing finds that consumption quality is a fundamental driver of consumers' satisfaction and well-being (see Oliver (1980) and subsequent literature for foundational contributions in this area). Sacrificing consumption quality might thus have an important negative effect on consumers' utility, which would not be captured fully by models of intertemporal consumption choice unless quality enters consumers' utility function. This possibility would be consistent with our evidence based on local retail market structures, whereby, ceteris paribus, consumers appear to maintain the average quality of their consumption in times of rising inflation if they have more choices across brands and retailers.

Absent direct survey-based data on consumers' reported satisfaction, we turn to an observable outcome that captures consumers' satisfaction with their living standards, that is, voting behavior. Anecdotally, experiencing high inflation during the 2021-2022 inflation surge has been proposed as an important driver of the shift from Biden to Trump in the 2024 US Presidential elections, which has been observed in a large fraction of US counties. Strikingly, we find that a county's inflation resilience is a strong predictor of this political shift, whereas local inflation rates per se are not after accounting for inflation resilience. In other words, it is not the mere exposure to rising inflation that explains voters' dissatisfaction but the extent

to which this exposure led local consumers to sacrifice their consumption quality through trading down. Counties with high inflation resilience, which largely kept their consumption quality constant, did not express their dissatisfaction against their living standards by voting against the incumbent administration in the 2024 US Presidential election. And, once we control for a county's inflation resilience, local inflation levels have no significant conditional association with voting outcomes.

This evidence sheds new light on the pros and cons of trading down in times of rising inflation for consumers and, incidentally, contributes to our understanding of how changes in consumers' economic choices transmit into choices in other realms, such as voting decisions, which in turn have major implications for the conduct of fiscal and monetary policy and macroeconomic outcomes.

Our work relates to several strands of literature. First, it contributes to research on household consumption under macroeconomic shocks. Jaimovich, Rebelo, and Wong (2019) document that during the Great Recession consumers traded down quality, with broad macroeconomic implications. We confirm a similar trading-down pattern during the recent inflation episode, but focus on its heterogeneity across space and sectors. Relatedly, Kaplan and Schulhofer-Wohl (2017); Jaravel (2024) study how inflation or price changes affect different income groups, and Argente and Lee (2021) examine how richer versus poorer households have different substitution opportunities. Our approach builds on this by introducing a direct measure of quality response to inflation and highlighting local cross-sectional differences.

Second, we connect to work on market structure and consumer choice under inflation. Cavallo and Kryvtsov (2024) shows that during the post-pandemic inflation surge, retailers' pricing strategies, e.g. frequency of discounts and "cheapflation" via product downsizing, affected the effective prices consumers pay. We complement these findings by showing the consumer side: where more alternatives exist, consumers can mitigate inflation's impact by swapping products, reducing the need for outright quality downgrades. In other words, competition and variety serve as cushions against inflation.

Third, our results on voting behavior ties into the literature in political economy and political science on what motivates agents to vote and what shapes satisfaction and dissatisfaction with incumbent parties in democratic regimes. We document a novel mechanism—trading down consumption quality in times of rising inflation—that links macroeconomic shocks to voting choices and hence feeds back into the conduct of fiscal policy and regulation.

Lastly, from a methodological standpoint, our paper shows the value of large-scale transaction data in measuring nuanced concepts like consumption quality, and even leverages new tools such as LLM-based brand rankings to validate and enrich traditional measures. By integrating these perspectives, our study offers a comprehensive analysis of how inflation affects households at the micro level and why it matters for both economic and social outcomes.

2 Data and Measurement of Consumption Quality

2.1 Data

To assess empirically the average reaction of US consumers' consumption quantity and quality in response to growing inflation and its heterogeneity across space and industries, our empirical analysis is based on data from several sources described below. For our main analyses, we aggregate the data at the US county-industry-month level.

Consumer Spending: Quantity and Quality

Our empirical inquiry requires data that allow us to measure consumption quantity and quality homogeneously for a large set of consumer transactions across space and industries. To this aim, we use novel data from *SafeGraph.Inc.*. The data bundle anonymized debit and credit card transaction information at the level of commercial points of interest (POIs) in the United States.[2] POIs are locations in which transaction information is recorded electronically. The data include not only transactions in stores in which a good or service is

[2]Note that this is not the POI-level data on foot traffic based on cellphone signal geolocation, which is also supplied by *SafeGraph.Inc.* and has been used in earlier research in economics and other fields.

physically transacted but also warehouses that process online orders, such as *Amazon Fulfillment Centers*. Including or excluding such warehouses, whose locations might not be close to the locations in which purchasers reside and whose share of total transactions has changed systematically during and after the COVID-19 pandemic, does not affect our results. The data exclude cash transactions, which are recorded at the brick-and-mortar store register level but are not processed through an electronic device.

Each POI is associated with a brand, which in turn belongs to an industry based on the North American Industry Classification System (NAICS). For instance, grocery stores are assigned the brand name of the retail chain or the independent business that operates them. We do not observe the brands of each specific good or service consumers purchase in a POI because the Universal Product Codes (UPCs) of each item purchased are not collected at the POI level. For instance, for NAICS3 445 (*Food and Beverage Retailers*), the data aggregate monthly transactions of brands such as *Walmart*, *Safeway*, *Whole Foods*, and all other chains or independent businesses in the industry at the level of POIs in which transactions are executed.

In November 2024, the data collected 207 million transactions made by 11 million consumers in 1,108,518 POIs covering 5,454 business brands and 21 NAICS3 industry codes across all 50 US states. *SafeGraph, Inc.* updates the data at the monthly frequency.[3] The sample we obtained spans from January 2020, which was before the COVID-19 pandemic diffused to the US, to December 2022. For each POI and month, we observe two sets of variables. The first set of variables includes transactions-based information: the total dollar amount spent in the POI, the total number of transactions, and the total number of unique customers that transacted in the POI over the month. The second set of variables consists of the POI-level information: POI's location (ZIP code, city, county, and state), brand name, and NAICS code. As we discuss in detail below, we use the information on the dollar value

[3]For the latest sample sizes and statistics, see *SafeGraph, Inc.*'s website: https://docs.safegraph.com/docs/spend-summary-statistics.

of transactions as well as on POIs' brands to compute our monthly measures of consumption quantity and quality by locations and industries.

To assuage the data's representativeness, Figure A.1 plots the cross-state correlation between the number of consumers whose transactions our data cover and the state-level population over the same sample period based on US Census data. The correlation value ranges from 0.8 to 0.9. In the Appendix, we provide additional formal tests to quantify the sample bias. Overall, we find that the sample is well-balanced demographically.

National and Local Inflation

We collect national and local measures of inflation from the Bureau of Labor Statistics. We use three measures of inflation for our analysis. First, we use year-on-year core Personal Consumption Expenditures (PCE) inflation at the monthly level as an economy-wide national inflation measure. Second, we construct a measure of local inflation that allows us to exploit cross-sectional spatial variation in inflation. We construct this measure iteratively: if available, we use monthly year-on-year Consumer Price Index (CPI) inflation measured at the metro area level. For cases in which CPI inflation at the metro area level is only available at the bi-monthly frequency, we linearly interpolate to impute the missing value between two available months;[4] if unavailable, we use monthly year-on-year CPI inflation measured at the US Census region level.

Demographics

Our demographic data come from the Census Bureau's American Community Survey. We consider the 1-year estimates, which we merge into our consumer spending data using FIPS codes. Because the 1-year ACS samples about 1% of the US population, it does not include observations in all US counties. We thus assign state-level averages of the demographic characteristics we observe to counties not covered in the years our sample spans. We collect the

[4]We use CPI inflation for regional measures because PCE inflation is not available at the metro area.

following variables: population, sex rate, median age, percent of African American residents, percent of White residents, percent of residents with a college degree, average family size, average household income, percent of homeowners with mortgages outstanding, percent of residents in the labor force, percent of residents 25-year-old or higher, median house value, median housing tax payments by residents.

2.2 Measuring Consumption Quality

The main challenge and objective of our analysis in terms of data is to construct homogeneous measurements of consumption quality at the geography-by-industry level using the transactions-based data in our sample.

Local Quality Measures — To measure consumption quality locally, we consider each POI as supplying a product, good or service, by a POI-specific brand.

In our baseline approach, we calculate, for each product, the average amount spent per transaction in each county in the first month of our sample. In this case, a higher-quality product corresponds to a product with higher average *per-transaction cost*. Therefore, our baseline measure of quality is based on the average cost of a product in a county before the start of the COVID-19 pandemic. In the baseline, we measure quality at the beginning of the sample period to avoid incorporating a look-ahead bias in our analysis, whereby consumption quality might depend on the changes in inflation and prices across brands, industries, and locations later in our sample period. As we describe below, we also compute alternative measures that relax this assumption for robustness purposes to find nearly identical empirical results. The baseline approach requires us to fix brands at the onset of the sample period and excludes new brands that enter the sample in later months.

Average per-transaction costs vary systematically across product types. For instance, durable goods such as cars have higher per-transaction costs than non-durable goods such as groceries. For this reason, in a second step, we categorize products' quality by comparing county-level average per-transaction costs within the same NAICS3 industry. Specifically, for

each county and NAICS3 industry, we split products into three equal-sized quality groups: *Top Quality* products are those in the top third of the distribution by per-transaction-cost within each county and NAICS3 industry. *Medium* and *Bottom Quality* products are those belonging to the middle and bottom third of the same distribution, respectively. Note that, based on this categorization, the same product/brand could be categorized as having different quality in different counties based on the distribution of brands available in different counties for the same NAICS3 industry. For instance, in NAICS3 445 (*Food and Beverage Retailers*), *Walmart* belongs to the bottom group of the quality distribution for the majority of US counties, *Whole Foods* to the top, and *Safeway* to the middle. Even though these rankings do not revert, *Safeway* belongs to the bottom or top of the quality distribution in some counties depending on the per-transaction costs of the other NAICS3 445 brands available in the same county.

Splitting products conditional on counties and industries aims to account for the location- and industry-specific factors that induce systematically different product costs. For example, the products in locations with higher purchasing power might have higher prices. Addressing this concern is important because many brands, and especially retail chains, engage in relatively uniform pricing across the US (DellaVigna and Gentzkow, 2019).

After categorizing each product/brand as belonging to the *Top*, *Medium*, or *Bottom* of the quality distribution by county and NAICS3, we compute the overall dollar value of the transactions executed in each county, NAICS3, and month in each partition of the quality distribution. Then, we compute the share of each partition on the overall dollar value of all transactions executed in that county, NAICS3, and month. In this way, we obtain our baseline *Top*, *Medium*, and *Bottom* quality consumption share at the county-NAICS3-month level, which we will use in the baseline empirical analysis. By construction, these three shares sum up to 1:

$$Consumption\ Share_{c,i}^{T/M/B} = \frac{\frac{\sum_{c,i} Dollar\ Value\ Transactions}{\sum_{c,i} Number\ Transactions} \times I_{c,i}^{T/M/B}}{\frac{\sum_{c,i} Dollar\ Value\ Transactions}{\sum_{c,i} Number\ Transactions}}, \tag{1}$$

where c indicates counties, i indicates NAICS3 industries, and $I_{c,i}^{T/M/B}$ is an indicator for the transactions that belong to the Top, Medium, or Bottom third of the distribution in county c and industry i in the first month of our sample.

To ensure that we can compare the consumption share across quality groups, for this local quality measure approach, we only keep in our sample county-by-industry units in which we observe at least three brands. Otherwise, we would be unable to assign each brand to a portion of the quality distribution or the assignment would be at random and create noise in the measure. Because several small and lowly populated counties have less than three brands within each NAICS3 industry, this restriction reduces the sample size non-trivially: our working sample for the analysis based on this local quality measure contains 1,162 unique brands across 323 US counties.

National Quality Measures — Our baseline approach results in substantial data loss by dropping counties with fewer than three brands within a given NAICS3 industry. If consumption behavior in the counties we drop is systematically different than in other counties, which ex-ante does not seem implausible, our baseline analysis might not be representative of the average US consumer's behavior. For this reason, we also construct alternative measures of product quality for robustness purposes, whereby we categorize brand quality at the US national level.

One advantage of defining quality tiers at the national level is that it avoids labeling a brand "high quality" in one county but "low quality" in another; the classification is uniform across locations. It also ensures that our measure is not unduly influenced by extreme local prices or limited local assortments. If a given county lacks any high-quality brands in an industry (for example, no luxury car dealerships), we assign a zero top-quality share for that category in that county, so that such counties remain in the analysis.

We acknowledge that price is an imperfect proxy for quality — some brands may command high prices due to factors other than quality, and quality perceptions involve attributes beyond cost. Therefore, we also construct an alternative national quality measure using an

independent qualitative assessment of brands. Leveraging recent advances in large language models that can handle a large amount of qualitative information homogeneously, we asked ChatGPT to rank the brands in our data on a five-point scale of "luxuriousness" based on general consumer perceptions. We use the term "luxurious" to give a clear directional gradient for the qualitative ordered categorization.[5] We aggregated those rankings into three groups (low, middle, high luxury) within each industry and mapped them to our data.

This AI-driven measure provides a check on our baseline categorization, incorporating non-price information such as brand prestige or customer reviews. We find that the ChatGPT-based rankings are strongly positively correlated with our price-based measure. As shown in Figure 1, which presents a binned scatter plot of the quality groupings assigned by ChatGPT against those determined by the average transactions cost in our national quality measure, higher average price is associated with a higher luxury rating on average in a cross-brand comparison. The slope of a binned scatter plot is approximately 20%, with a t-stat over 10. At the same time, the imperfect fit ($R^2 = 0.04$) indicates that the two measures capture somewhat different nuances. We use this alternative classification in robustness tests and find that our main results are qualitatively unchanged, which speaks to the external validity of our baseline quality measure.

Other Measures — Our baseline location-based measure of quality also faces the potential concern that in certain counties, products in the same industry might be highly homogeneous but by construction they would be assigned to different quality groups even if their costs are similar. For the sake of robustness, we thus also replicate our results when only considering counties and industries where the difference between the average transaction cost of the top- and bottom-quality product groups is above the upper tercile of the full-sample distribution. This restriction, by construction, estimates the results on only one third of the baseline sample.

[5]The prompt we use is *"Here is a list of brands in the US. Could you assign them into five groups based on how luxurious they are?"*

The other restriction to our quality measure we relax in robustness analysis is basing it on the average cost per transaction at the beginning of the sampling period. As discussed above, this restriction is important to avoid a look-ahead bias in our measure. At the same time, our measure based on a single month of underlying data might be noisier than a measure based on the full dataset. We thus assess the robustness of our results when relaxing this restriction and constructing our county-by-industry local quality measure based on the average cost per transaction by brands throughout the sample period rather than only in the first month.

2.3 Summary Statistics

Table 1 provides summary statistics of our main working sample at the county-NAICS3-month level using the baseline local consumption quality measure. Overall, the sample includes 44,686 units of observation.

In terms of consumption variables, we can see that average spending is right-skewed. On average, consumers spend $76.4 on each product transaction but the median is much smaller (about $36) and the standard deviation is almost twice as large as the mean (about $130). Unsurprisingly given our consumption quality categorization based on average brand's cost per transaction, the average amount of spending increases with quality. Consumers spend on average $42.2, $73.3, and $138.5 dollars on low-quality, medium-quality, and high-quality products, respectively, for each transaction. All these measures are also right skewed, that is, the overall skewness of the average cost per transaction is not driven by the top quality category but is a common feature of all quality groups.

For each product in a given county and given year-month, our sample includes, on average, about 267 transactions made by an average of about 154 consumers.

Moving on to measures of inflation, over our sample period, the average national PCE inflation measured at the level of unit of analysis (the county-NAICS3-month) is 3.42% and the average national CPI inflation is 4.69%. PCE inflation is lower in both average magnitude and standard deviation across units (1.73 ppt versus 2.98 ppt, respectively). Local and

national CPI inflation display very similar distributions. The average and standard deviation of Local CPI inflation are 4.68% and 3.10%, respectively.

3 Inflation Resilience

Our analysis proceeds in two stages. First, we estimate the average relationship between inflation and consumption quality. Second, we move on to assess the heterogeneity of this reaction across space and industry and investigate the drivers of this heterogeneity.

3.1 Estimating Inflation Resilience: Baseline Specification

Our baseline specification regresses the quality share on measures of inflation, including a rich set of controls. Specifically, we estimate equations of the form:

$$Y_{c,i,t} = \beta \cdot \pi_{c,t} + \boldsymbol{X}_{c,i,t}\Gamma + \alpha^{FE} + \epsilon_{c,i,t}, \tag{2}$$

where $\pi_{c,t}$ is the inflation rate in county c and year-month t or a national inflation measure common to all counties interacted with local traits. $\boldsymbol{X}_{c,i,t}$ represents demographic controls such as county-level income, population, age, education, race, and other variables in Section 2.1. We include county fixed effects α_c to account for time-invariant differences across counties such as baseline consumption preferences or persistent cost-of-living differences, and industry fixed effects α_i to capture any systematic differences across industries. Time fixed effects α_t, the year-month dummies, absorb aggregate shocks such as national inflation or pandemic effects common to all counties. In some specifications we replace α_c with more granular effects, like county-by-industry fixed effects $\alpha_{c,i}$, to ensure we exploit within-location-and-industry variation.

The dependent variable $Y_{c,i,t}$ is either the consumption share of bottom-quality, medium-quality, and top-quality products $T/M/B_{c,i,t}$ or the difference of consumption shares across quality tiers, $(M-B)_{c,i,t}$ or $(T-B)_{c,i,t}$. Specifically, when we study the difference of con-

sumption shares between top-quality and bottom-quality products, the coefficient $\beta_{(T-B)}$ measures the average *inflation resilience*. The null hypothesis $\beta = 0$ indicates that inflation has no association with local consumption adjustments along the quality dimension. A negative β indicates trading down on average by local consumers. We cluster standard errors at the state-by-year-month level. Table 2 presents the estimates. Columns 1-3 consider the local bottom-quality, medium-quality, and top-quality consumption shares separately as dependent variables. Given that the shares of these three groups sum to one, the sum of the three coefficient estimates must net out to zero. Column 1 shows that, In a simple baseline specification without fixed effects or controls, each ppt increase in regional inflation is associated with a 0.14 ppt increase in the low-quality local consumption share. This increase is balanced by drops of similar sizes in the top-quality (0.07 ppt) and medium-quality (0.06 ppt) consumption shares.

Column 4 focuses on the differential between top- and bottom-quality consumption shares. The point estimates, which by construction are consistent with the estimates for each share separately, show that the relationship between local inflation and the overall redistribution of consumption across top- and low-quality groups is negative and statistically significant at the 1% level. After adding county fixed effects, industry-by-year fixed effects, and demographic controls, the coefficient stabilizes around -0.68 and remains statistically significant. This implies that a one ppt increase in inflation results in a 0.68 ppt drop in local top-quality consumption shares relative to low-quality shares. In column 7, we also introduce a set of interaction terms between state fixed effects and national-level inflation, which allow us to absorb the reaction of time-invariant local characteristics to the same inflation shocks and only allows for within-state local variation in inflation to identify our coefficient estimates. Our estimates do not change directionally, statistically, or in terms of magnitude.

Overall, the baseline results confirm that, on average, higher inflation leads consumers to shift towards lower-quality consumption, which is consistent with evidence of inflation-driven "trading down" consumption behavior.

3.2 Intensive and Extensive Margins of Quality Adjustment

Next, we explore how this adjustment across quality categories happens. On the one hand, consumers may adjust along the extensive margin, whereby they engage in more transactions of lower-quality goods and services as inflation grows. On the other hand, consumers might not transact more in low-quality goods but if the price per transaction of lower-quality goods grew more than that for higher-quality goods, that is, lower-quality goods' prices reacted more in times of growing general inflation, the bottom-quality consumption share would increase mechanically (intensive margin). Assessing the empirical relevance of these non-mutually-exclusive margins is important because each margin points to different potential drivers of inflation resilience.

We assess the margins of adjustment using a version of the baseline specification whereby we replace consumption shares with proxies for the intensive and extensive margins of adjustment.

First, we consider the intensive margin of adjustment, measured by the logarithm of the cost per transaction for goods of different qualities. Panel A of Table 3 shows that, unsurprisingly, higher inflation correlates with a higher cost per transaction across all quality groups, though the cost per transaction of goods and services at the bottom-quality group (column 1) increases more than those at the top-quality group (column 3), 2.59 version 1.74, both significant at the 1% level. The difference is -0.86, as shown in column 4. After adding demographic controls and restricting the variation within industries and counties (column 5-6), however, this difference becomes insignificant, and even switch signs to positive. Compared to the results in Table 2, the joint findings suggest that the intensive margin, that is, the differential reactions of the prices of goods and services across the quality distribution in face of rising inflation, are unable to explain the variation in inflation resilience we have documented earlier.

We then move on to consider the extensive margin. Our first measure of the extensive margin is the average number of transactions per consumer by quality tier in each county,

industry, and year-month. Panel B of Table 3 shows that the number of transactions per consumer, if anything, declines for bottom-quality goods and services and increases for others. However, these patterns are due to systematic differences in the composition of industries and county-level time-invariant characteristics as well as economy-wide shocks. Once we partial out these characteristics to isolate the role of local inflation, we find that the relative number of top-quality transactions over bottom-quality transactions declines as inflation increases. The second measure of the extensive margin is the logarithm of the total number of consumers actively purchasing in each quality tier. This variable captures the extent to which consumers stop transacting in a quality group and/or start transacting in another quality group altogether. Panel C of Table 3 shows that this component of the extensive margin of adjustment is important to explain the cross-sectional variation in inflation resilience. Indeed, even though the total number of consumers declines across the board (columns 1 to 3), which indicates that consumers concentrate their consumption within fewer quality groups rather than splitting it across all quality groups, we find that it declines substantially more for higher-quality than lower-quality groups. The magnitude of this association is larger once we restrict the variation within industries, counties, and accounting for time-varying shocks.

Overall, our analysis of the margins of adjustment provides direct evidence that the baseline estimates of inflation resilience capture an active choice by households in terms of transacting and purchasing behavior rather than a mechanical effect due to differential relative price changes across goods of different quality.

4 Heterogeneity of Inflation Resilience Across Space and Industries

We now turn to the main aim of our paper, that is, assessing the cross-sectional heterogeneity of inflation resilience across space and industries as well as its drivers. We have already noted the wide dispersion in estimated resilience (the sensitivity of quality share to inflation) across

counties. In this section, we delve deeper into which types of counties (and industries) are more or less resilient, and why.

To do this, we adjust the baseline specification to include interactions of local inflation with county fixed effects and industry fixed effects, which allows us to estimate inflation resilience for each of 323 counties and 21 industries in our data:

$$(T-B)_{c,i,t} = \sum_{c} \beta_c \cdot \eta_c \cdot \pi_{c,t} + \sum_{i} \beta_i \cdot \eta_i \cdot \pi_{c,t} + \boldsymbol{X}_{c,j,t}\Gamma + \epsilon_{c,i,t}, \tag{3}$$

where η_c and η_i denote the full set of fixed effects at the county and industry levels.

Estimating equation (3), we find that inflation resilience varies markedly across space. The standard deviation of our county-level estimates of inflation resilience $\hat{\beta}_c$ is 1.61 ppt, with an interquartile range (IQR) of 1.71 ppt. These values are nearly three times as large as the average inflation resilience in the baseline specification, which is -0.68 ppt in column 6 of Table 2 and -0.50 ppt in column 7 with the inclusion of state-by-national inflation interactions in the specification. This substantial across-county variation indicates that locations exist in which inflation resilience is strongly negative, that is, local consumers trade down dramatically in the face of growing inflation, while at the same time consumers in other locations do not trade down at all and keep consuming at the top of the quality distribution despite growing inflation, to the extent that the share of overall expenditures in goods and services at the top of the quality distribution ends up being higher than that at the bottom of the quality distribution. The rest of our paper aims at describing this variation and investigating its drivers.

We start by visualizing the spatial distribution of inflation resilience across U.S. counties in our sample in Figure 2. Counties are categorized into four equal-sized groups based on their estimated inflation resilience. Red triangles denote counties in the bottom quartile (lowest resilience), where consumers exhibit the highest propensity to trade down in response to inflation. Orange circles and light-blue diamonds represent counties in the second and third

quartiles, respectively, while dark navy squares indicate counties in the top quartile with the highest inflation resilience.

Figure 2 shows that, first of all, our sample includes US counties scattered across the US with a sizable representation of all US Census Regions and most states. Generally, we observe no obvious spatial patterns in inflation resilience; counties with different levels are present in most regions. However, some concentrations are apparent: southern Florida and the Mid-Atlantic coast (around New York City and New Jersey) show a cluster of low-resilience counties. Conversely, Southern Texas predominantly features counties with moderate resilience. The central US Midwest contains very few counties in the top resilience quarter, although values vary across the other quartiles.

Our estimates of inflation resilience at the industry level $\hat{\beta}_i$ also suggest a substantial variation. We find that the standard deviation is 1.41 ppt and the IQR 0.79 ppt, both higher than the average inflation resilience in the baseline specification. Specifically, we find that inflation resilience is lowest, that is, trading down behavior is most pronounced, for the consumption of services in *Child Care and Family Services* (NAICS 624) and *Amusement and Recreation Services* (NAICS 713). In contrast, inflation resilience is highest in *Electronics and Appliance Retail*. For the consumption of non-durable goods, such as those purchased in grocery stores, which is captured by *Food and Beverage Retail* (NAICS 445), we estimate an inflation resilience of -1.04, which is only slightly lower than the average resilience in our sample.

This industry distribution of inflation resilience suggests that, in face of rising inflation, consumers seem to be trading down substantially in the consumption of non-durable goods and services, and especially services that might not be interpreted as necessities. On the other hand, they seem to not be trading down at all in the consumption of durable goods. These patterns open a set of interesting questions, such as what features of goods and services lead consumers to adjust the quality of their consumption in times of growing inflation. For

instance, whether goods and services are perceived as necessities or conspicuous consumption might lead consumers to react differently.

4.1 Drivers of Inflation Resilience 1: Income and Pre-Inflation Consumption Quality

The substantial spatial and industry variation in inflation resilience prompts us to dig deeper into underlying what observables help explain this variation in how consumers adjust their consumption quality and quantity in response to inflation.

We first consider demographic characteristics. Theory and existing evidence on trading down suggest that, because high-income households generally consume higher-quality goods, they have more scope to reduce consumption quality keeping their quantity unchanged; by contrast, lower-income households, who already spend mostly on lower-quality goods, have less scope to trade down (Argente and Lee, 2021).

To assess this channel, we adjust our baseline specification as follows:

$$(T-B)_{c,i,t} = \beta \cdot \pi_{c,t} + \beta_H \cdot High\ \ Income_c \cdot \pi_{c,t} + \beta_M \cdot Medium\ \ Income_c \cdot \pi_{c,t} + \boldsymbol{X}_{c,j,t}\Gamma + \epsilon_{c,i,t}, \quad (4)$$

where, $High\ Income_c$ and $Medium\ Income_c$ are indicators for county c, taking a value of one if the county falls in the top or middle third of the median household income distribution, respectively, and zero otherwise. The coefficients β_H and β_M represent the additional association between inflation and the difference in top-quality versus bottom-quality consumption shares for counties in the top and middle thirds of the income distribution. The coefficient β captures this association for the omitted category, the bottom third of the county-level median income distribution. This specification enables us to evaluate whether income levels exhibit a linear or non-linear relationship with inflation resilience across three segments of the income distribution.

The estimates provide a clear pattern: higher-income counties exhibit less inflation resilience, meaning they experience larger declines in consumption quality for a given inflation shock. Panel A of Figure 3 illustrates this point by plotting the distribution of inflation resilience estimates for counties in which the median household income is equal or greater than \$100K (green solid line) versus those equal or lower than \$50K (orange dashed line). The distribution of inflation resilience for high-income counties is shifted to the left (more negative), indicating a stronger trading-down response, whereas low-income counties cluster more around zero or mild negatives, suggesting smaller quality adjustments.

In column 1 of Table 4, we formally compare high- vs. low-income counties. We find that inflation resilience is 0.22 ppt lower in high-income counties, which is approximately 47% lower than that in low-income counties. This role of income does not appear to be linear in our data given that counties in the middle of the distribution do not behave economically or statistically differently from lower-income counties.

To directly test whether heterogeneous inflation resilience by income is due to the differential scope of trading down for consumers along the income distribution, we construct a variable measuring each county's pre-inflation-surge low-quality consumption share. This serves as a proxy for how much trading-down space consumers in that county had. And, column 2 of Table 4 confirms the mechanisms: counties in the top quartile of the initial bottom-quality share distribution show a significantly smaller response to inflation (i.e. a positive interaction term, indicating higher resilience) compared to counties in the bottom quartile of initial bottom-quality share.[6] In fact, after controlling for this initial basket composition, the direct effect of income becomes much weaker and statistically insignificant in most specifications.

To further corroborate the importance of this channel, we perform a falsification exercise by considering spatial variation in other demographics: education levels and age distributions.

[6]Panel B of Figure 3 confirms that the distribution of inflation resilience is shifted to the left for counties in the top quartile of the initial bottom-quality share distribution,which has deeper scope for trading down by local consumers.

Counties that differ systematically in terms of their consumers' education levels or age might display systematic differences in the average consumption baskets, but not necessarily in the distribution of quality of such baskets above and beyond what is explained by income levels. Instead, if demographic variation mattered through channels that differ from local households' ability to trade down based on pre-inflation-surge consumption quality shares, we would find significant associations.

Columns 3-4 of Table 4 show that we find no systematic relationship with inflation resilience once we account for income and the initial consumption quality shares.

4.2 Drivers of Inflation Resilience 2: Local Retail Market Structure

Figure 3 shows that within each group of counties sorted by income levels or the pre-inflation bottom-quality consumption share, there still exists substantial cross-sectional variations in inflation resilience. This is prima facie evidence that dimensions not captured by income must contribute to explaining the spatial variation in inflation resilience. Moreover, the results in section 3.2 on the margins of consumption quality adjustment indicate that potential drivers should relate to consumers' ability to migrate across consumption quality groups and to actively change their transaction habits as inflation grows rather than passively facing the differential reaction of prices across different quality groups.

In this subsection, we investigate how the characteristics of local product markets might influence inflation resilience. As hypothesized in the introduction, the availability of a broad range of products and brands in a locality could cushion the impact of inflation on consumption quality because consumers could find similar-quality products whose prices react less than those they typically consume and hence shield their consumption baskets from inflation by maintaining their consumption quality unaltered.

We measure local brand variety in two ways: (i) the total number of distinct brands that consumers in a county purchase across all industries, and (ii) the number of brands within

each county-industry bin.[7] Local brand variety is a sticky county characteristic that cannot easily adjust over a short period due to adjustment frictions such as adapting distribution networks, obtaining licenses, and building consumer awareness through advertising. We consider local brand variety because lower-quality brands tend to be more concentrated than higher-quality ones. For example, in grocery retail, a single lower-quality brand like Walmart captures a large share of consumers and transactions, while higher-quality brands are typically smaller and more geographically dispersed. Higher brand variety offers higher-quality consumers more alternatives to switch among similar brands as inflation rises, potentially allowing them to maintain consumption quality and quantity. Conversely, in areas with lower brand variety, fewer high-quality options exist, potentially forcing consumers to shift to lower-quality alternatives to avoid increasing spending per unit or reducing quantity.

We test this channel in Table 5. Columns 1â€“2, show that counties in the top variety group are more resilient to rising inflation, with a markedly smaller drop in quality share, compared to those in the low variety group. For instance, moving from a county at the 50th percentile of brand count to one at the 90th percentile is associated with roughly a one-third higher inflation resilience (i.e., the quality drop is 37% lower). When looking within industries (column 2), an industry in a location with high brand variety — many brands are available in that local market— shows up to a 75% higher resilience than the same industry in a low-variety location. Panel C of Figure 3 visualizes this point by plotting the average inflation-quality slope in markets with a low (solid green line) versus a high (dashed orange line) number of available brands. The high-variety markets have a slope much closer to zero (flatter), indicating little quality reduction, whereas low-variety markets have a steep negative slope.

The second feature of local retail market structure our data allow us to examine is retail concentration, which captures potential frictions affecting consumers' access to high-quality brands and retailers beyond the mere number of brands available. Intuitively, in locations

[7]In this experiment, we consider two-digit NAICS industries instead of the finer three-digit NAICS industry partitions. This is because many county-by-NAICS3 pairs include a single brand.

with fragmented retail markets where multiple retailers sell goods from the same brands, consumers have the possibility to shift across retailers that use different pricing strategies in times of inflation to keep consuming in the same quality tier by chasing lower price increases. By contrast, in locations with high sales concentration where a few large retailers dominate transactions and impose uniform pricing strategies (e.g., see DellaVigna and Gentzkow (2019)), consumers should find it harder to maintain the quality of the brands they consume without facing higher prices. Consistently, we find that inflation resilience is economically and statistically lower in counties in which retail transactions are more concentrated relative to those in which transactions are more dispersed across retailers.

We use POI transaction data to quantify this aspect of the local retail structure by computing the Herfindahl-type concentration of transactions across local retailers in each county. A higher concentration indicates the dominance of large retailers, whereas lower concentration suggests a greater presence of medium- and small-sized retailers.

We then ask whether counties with more concentrated retail markets exhibit different inflation resilience. The data show that they do: high-concentration counties have significantly larger drops in quality share in response to inflation than low-concentration counties. We illustrate this point first in the raw data in Panel D of Figure 3: the relationship between inflation and quality share is much more negative (steeper) in high-concentration locales than in low-concentration locales. This result ties into the idea of search frictions and awareness: when a few chains dominate, consumers might be less aware of niche alternatives or face higher search costs to find them, leading to more pronounced trading down within what they know.

Columns 3-4 of Table 5 present results from estimating the interactive specification of Regression (4). Interaction groups are based on the distribution of the Herfindahl-Hirschman Index (HHI) of transactions, computed at the county level in column 3 or at the county-by-industry level in column 4. In both specifications, we find that inflation resilience is lower in areas with higher transaction concentration, consistent with our conjecture. Specifically,

in counties at the top of the HHI distribution (column 3), resilience is about 47% lower compared to those at the bottom. Furthermore, column 4 shows that in low-concentration groups, consumers do not significantly shift consumption quality as inflation rises. However, in higher-concentration groups, they react substantially to increasing inflation, a finding consistent with the results for brand variety in column 2.

It is worth noting that brand variety and retail concentration are themselves correlated with income — wealthier, more urban counties tend to have both more variety and possibly more concentration, if dominated by major chains. In our analyses, we control directly for income levels. This implies that even for two areas with similar income, the one with a richer set of consumer options will handle inflation better in terms of sustaining consumption quality. These findings contribute a novel perspective to our understanding of the interplay of inflation with local retail market organization. In policy terms, it suggests that promoting competition and entry in retail markets could make communities more resilient to inflation by giving consumers alternative ways to adjust their consumption in the face of rising inflation above and beyond trading down and reducing consumption quality.

5 Endogeneity Concerns and Robustness

5.1 Instrumental Variable Strategy

The larger estimates of inflation resilience in our most restrictive specifications that include a large set of controls and fixed effects suggest that potentially unobserved characteristics, if anything, might bias our baseline OLS specifications downwards. Still, one might be worried that our evidence so far is subject to a set of endogeneity concerns. For instance, regions with higher productivity often experience increased local aggregate supply, which can lead to elevated inflation levels. Simultaneously, heightened productivity may raise real wages, prompting households to allocate more spending towards higher-quality goods. Consequently,

this dual effect may lead to an upward bias in the estimated coefficient for inflation resilience, β.

To tackle these concerns more directly, we propose an instrumental-variable strategy based on a shift-share argument. Specifically, we exploit two ingredients: an external cost shock due to price increases in manufacturing good inputs imported from China over the period 2021−2022 ("shift") and the pre-2021 exposure of each US state's economy to manufacturing ("share"). The idea is that input price increases imposed outside the US, which are largely driven by supply conditions in China, led to higher production costs and inflationary pressure in US locations that had a larger manufacturing sector and hence were more exposed to such external shocks, irrespective of local demand conditions.

To implement this strategy, for each state s, we construct the pre-inflation-surge manufacturing share, MS_s, as the ratio of value added from the manufacturing industry to the total state-level GDP, averaged over the period between 2015 to 2019. Then, for each year-month t, we get the year-on-year price inflation on the goods that US manufacturing industry imports from China, IC_t. IC_t is therefore a proxy for the general costs of input from China. We use China because it is the country from which the US has the largest share of imports right before our sampling period. The instrument variable is then constructed as

$$Instrument_{s,t} = IC_t \times MS_s. \tag{5}$$

For this measure, IC_t is the time-varying cost shifter, and MS_s is the state-level pre-sample period exposure to the cost shifter. The assumption on relevance is that when the intermediate goods that the manufacturing industry imports from China get more expensive, the overall price level will be higher, and this is strong for places with an ex-ante larger share of value added from the manufacturing industry. Meanwhile, controlling for location fixed effects, goods inflation in China does not affect the difference in the consumption share between high-quality and low-quality products other than through changing inflation.

With the measure $Instrument_{i,t}$, we fit the following two-stage-least-square estimation (2SLS):

$$\pi_{c,t} = b_1 \times Instrument_{s,t} + \boldsymbol{X}_{c,i,t}\Gamma_1 + \epsilon_{1,c,i,t},$$
$$(T-B)_{c,i,t} = \beta_{IV} \cdot \widehat{\pi}_{c,t} + \boldsymbol{X}_{c,j,t}\Gamma_2 + \epsilon_{2,c,i,t}. \tag{6}$$

β_{IV} measures the effects of inflation induced by the regional heterogeneous exposure to the goods inflation in China on the difference between top-quality and bottom-quality consumption shares.

The first-stage results, as shown in column 1 of Table 6, confirm that this China-input inflation instrument strongly predicts higher local inflation. A one-standard-deviation increase in $Instrument_{i,t}$ corresponds to a 0.28 ppt rise in local inflation. This result is consistent with the idea that the higher prices of goods imported by U.S. manufacturers from China disproportionately affect regions with greater initial exposure to manufacturing industries. Importantly, the inclusion of county fixed effects ensures that time-invariant heterogeneity across locations, that is, characteristics like local institutions, local regulation, local economic activity structure, or time-invariant local tastes and preferences do not drive this variation. The significant elasticity and the high value of F-statistic 78.84 indicate no weak instrument problem.

This instrument would satisfy the exclusion restriction as long as Chinese input cost shocks did not affect local households' consumption quality choices through direct channels other than local inflation. Although this assumption is untestable, we believe it is economically plausible. One could imagine, for example, that areas with big manufacturing sectors might also see employment changes from global shocks; however, our period of analysis is short and dominated by price movements rather than quantities, and we control quite flexibly for direct demand effects with fixed effects and income controls.

We re-estimate the inflation effect on quality using our 2SLS specification. Across various specifications, the IV coefficient on inflation is significantly negative and larger in magnitude than the OLS estimates in Table 2. For example, in the specification analogous to specification (6) in Table 2, the IV point estimate suggests that a 1 ppt increase in inflation leads to roughly a 1.0 ppt decline in high-quality share, roughly 50% more negative than the OLS coefficient, 0.68 ppt. This fact is consistent with OLS being biased towards zero by attenuation or simultaneous shocks that muted observed inflation. This gives us confidence that the relationship we document is indeed a causal effect of inflation on consumption behavior.

5.2 Robustness: Alternative Quality Measures and Subsample Analysis

In this section, we conduct a set of robustness checks to ensure our findings are not artifacts of specific definitions or sample selection.

First, we re-estimate equations (2) and (6) using the national quality measures. The results, presented in Panel A of Table 7, again demonstrate a negative relationship between regional inflation and the difference in consumption shares of top- versus bottom-quality products. Due to the inclusion of numerous zero consumption shares for county-industry pairs with fewer than three products, the variability of the dependent variables is reduced. Consequently, the estimated magnitudes of inflation resilience are also diminished.

Second, we verify that using the ChatGPT-based quality measure, which relies on qualitative brand reputation rather than price, yields similar patterns. Despite the imperfect correlation between the cost-based measure and ChatGPT's classifications, we continue to find a significant trading-down effect of inflation and similar heterogeneity across income and variety dimensions, as shown in Panel B of Table 7. Column 5 indicates that each percentage point increase in inflation is associated with a 0.31 percentage point larger reduction in the consumption share of top-quality products compared to low-quality product.

Third, we restrict the sample to instances where the separation between high- and low-quality tiers is stark, for example, industries where top vs. bottom average prices differ by a large margin. Our baseline could be noisy in industries where "high" vs "low" quality are not distinct and hence products are assigned to different quality groups even though they have minimally different prices. In Panel A of Table 8, we refit regression (2) while focusing on the counties and industries where the quality differential between the top- and bottom-quality groups are above the upper tercile. In particular, when focusing on the counties and industries with clearly differentiated quality tiers, a one percentage point increase in inflation results in a 1.25 ppt larger quality gap, suggesting our main estimates are, if at all, underestimating the true effect.

Fourth, we modify the quality measure from using the average cost per transaction at the beginning of the sampling period to those based on the average cost throughout the sample period. The former restriction is to avoid a look-ahead bias but it may contains more noises with only single month data. Relaxing this restriction and using the full-sample transactions, we find that the findings are qualitatively the same, but the magnitudes are much larger. Meanwhile, while there is no significant adjustment in the gap of average cost per transaction or average transaction per consumer, 4.31% more consumers opt for low-quality products relative to high-quality ones.

6 Inflation Resilience and Voting Behavior

In the last part of the paper, we aim to shed light on whether inflation resilience is perceived as a negative or positive outcome by consumers. This question is inspired by our evidence, which suggests a trade-off. On the one hand, and consistent with the standard interpretation of trading-down behavior in macroeconomics, the possibility of trading down is an advantage for consumers, who can shield their bundles and consumption quantities from the negative effects of inflation by adjusting along the quality margin. Under this interpretation, for

instance, the fact that higher-income consumers can trade down while low-income consumers cannot is interpreted as an advantage for high-income consumers in times of rising inflation.

At the same time, though, a large literature in social psychology and marketing finds that consumption quality is a fundamental driver of consumers' satisfaction and well-being (see Oliver (1980) and subsequent literature for foundational contributions in this area). Sacrificing consumption quality might thus have an important negative effect on consumers' utility, which would not be captured fully by models of intertemporal consumption choice unless quality enters consumers' utility function. This possibility would be consistent with our evidence based on local retail market structures, whereby, ceteris paribus, consumers appear to maintain the average quality of their consumption in times of rising inflation if they have more choices across brands and retailers.

Absent direct survey-based data on consumers' reported satisfaction, we turn to an observable outcome that captures consumers' satisfaction with their living standards, that is, voting behavior. Anecdotally, experiencing high inflation during the 2021-2022 inflation surge has been proposed as an important driver of the shift from Biden to Trump in the 2024 US Presidential elections, which has been observed in a large fraction of US counties.

We use county-level data on the change in the Democratic vote share from the 2020 Presidential election (when inflation was low) to the 2024 election (after the inflation surge).[8] We relate this change to two key variables: the county's cumulative inflation rate over the period and our measured inflation resilience based on 2021−2022 consumption behavior. We estimate the following regression:

$$\Delta Democratic\ Share_{c,2020\to2024} = \delta \cdot Inflation\ Resilience_c + \beta \cdot \bar{\pi}_c + \boldsymbol{X}_c\Gamma + \epsilon_{c,t}, \quad (7)$$

where $Democratic\ Share_{c,2020\to2024}$ is the difference between the Democratic candidate's vote share in 2024 and in 2020 Presidential Election in county c, $Inflation\ Resilience_c$ are the standardized county-level estimates of inflation resilience, $\bar{\pi}_c$ is the time-series average of local

[8]The vote share data is collected from the *Wall Street Journal*'s Presidential Election Live Results.

inflation at the county level over our sample period, and X_c is the average of the demographic controls at the county level over our sample period. Note that, Since the unit of observation is a county, this specification does not include the set of fixed effects in the most restrictive version of our analyses run at the county-industry-month level earlier in the paper.

Table 9 shows that without controls, as shown in column 1, a standard deviation lower resilience correlates with about a 2 ppt larger decrease in Democratic vote share. In column 2, adding the local inflation rate, the coefficient on resilience remains almost unchanged, while the coefficient on local inflation rate is near zero and insignificant. Adding demographics controls such as income and education in column 3, the resilience effect persists, suggesting it is not simply capturing that poorer or otherwise different counties both traded down more and swung votes. In contrast, the level of inflation to which voters were exposed has no economic or statistical correlation with the change in voting choices. In fact, one might expect poorer counties to have lower resilience (as we found they had higher resilience due to less scope to trade down), so if anything, controlling for income might even strengthen the resilience-vote relationship.

Figure 4 visualize this finding. Panel A shows a bin-scatter plot of the 2024 Presidential Election county-level vote share for the candidate of the incumbent Democratic party, Kamala Harris, against our county-level inflation resilience estimates. We detect an evident positive correlation between the two variables. Counties in which inflation resilience during the 2021 inflation surge was lower or stronger trading down behavior were substantially less likely to vote for Kamala Harris than counties in which inflation resilience was higher. Panel B shows that inflation resilience correlates positively also with the change in the county-level vote share of the 2024 Democratic candidate (Kamala Harris) relative to the county-level share of the 2020 Democratic candidate (Joe Biden). This change is a crucial outcome because it explains the extent to which the average local household's voting choices changed, which ultimately led to the demise of the incumbent administration and the election of the Republican candidate, Donald Trump.

Our findings imply that voters responded to how inflation impacted their consumption quality rather than just the fact that prices were higher. These findings propose a new perspective into our understanding of the consequences of trading down behavior, which does help households to shield their consumption quantities from the effects of inflation but creating by-products, such as dissatisfaction and disutility, that standard intertemporal consumption models might overlook.

This section underlines another relevant point: inflation resilience has real consequences beyond consumption choices; it also affects other fundamental choices consumers make. It can potentially affect how economic shocks translate into social outcomes like political stability and trust. In terms of voting, contrary to the most common views, exposure to inflation *per se* does not appear to be an important driver of US voters' switch from supporting Joe Biden in 2020 to supporting Donald Trump in 2024. Rather, the extent to which consumers changed the quality of their consumption in response to rising inflation appears as a key driver of this switch. Policymakers often focus on headline inflation statistics, but our evidence suggests they should also pay attention to metrics of household adaptation and strain, as these could be leading indicators of socio-political feedback effects.

7 Conclusions

Using large-scale transaction data, we define and measure inflation resilience across U.S. localities and industries. We find that, on average, consumers trade down quality during inflation spikes, but that this effect varies widely. Higher-income regions and markets with limited variety were much more prone to downgrade quality, whereas poorer areas and rich, competitive markets were more buffered. These patterns hold under rich controls and even under an instrumental-variable strategy that exploits exogenous supply shocks. Crucially, we show that this differential resilience has real consequences beyond consumption: it helps

predict political swings in a major election, suggesting that welfare losses from inflation are felt and acted upon in communities.

Our study contributes to understanding how inflation affects welfare in a nuanced way. It suggests that policy debates on inflation should not only focus on headline rates but also on how local market structures and consumer options mediate the burden. Measures that preserve or enhance market competition and consumer choice could increase resilience and thus welfare during inflationary periods. Future research could explore interventions such as subsidies for staple goods or support for small retailers that target resilience.

In conclusion, by quantifying inflation resilience, we reveal that it is not inflation *per se* but the erosion of consumption quality that truly hurts people's living standards and their political preferences. Understanding and enhancing resilience could make the economy more robust to future shocks.

References

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

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

DellaVigna, S. and M. Gentzkow (2019, 06). Uniform pricing in u.s. retail chains*. *The Quarterly Journal of Economics 134*(4), 2011–2084.

Jaimovich, N., S. Rebelo, and A. Wong (2019). Trading down and the business cycle. *Journal of Monetary Economics 102*, 96–121.

Jaravel, X. (2021). Inflation inequality: Measurement, causes, and policy implications. *Annual Review of Economics 13*(1), 599–629.

Jaravel, X. (2024). Distributional consumer price indices. Technical report, LSE Working Paper.

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

Oliver, R. L. (1980). A cognitive model of the antecedents and consequences of satisfaction decisions. *Journal of marketing research 17*(4), 460–469.

Figure 1: Correlation of Consumption Quality Measures: Average Cost per Transaction by Brand vs. AI-based Brand Categorization

This figure gives the binned scatter plot of quality as measured by the average cost per transaction and that labeled by ChatGPT-4o.

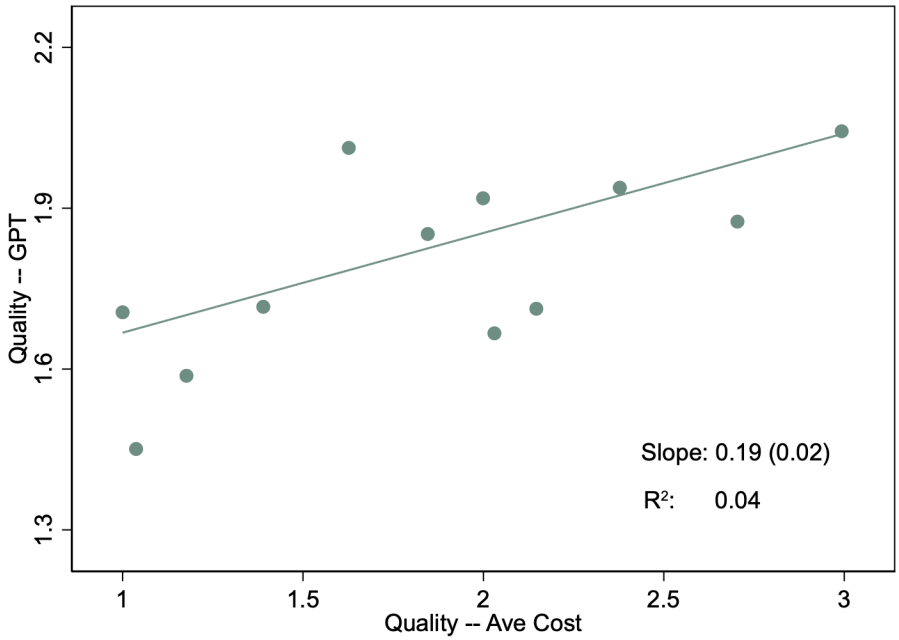


Figure 2: Spatial Distribution of Inflation Resilience Across US Counties

This figure depicts the spatial distribution of US county-level inflation resilience during the 2021-2022 inflation surge in longitude-latitude space. Each symbol refers to a US county for which Safegraph collects POI-based transaction data. Red triangles, orange diamonds, light blue circle, and navy squares are counties in the first (bottom), second, third, and fourth (top) quarters of the distribution by inflation resilience.

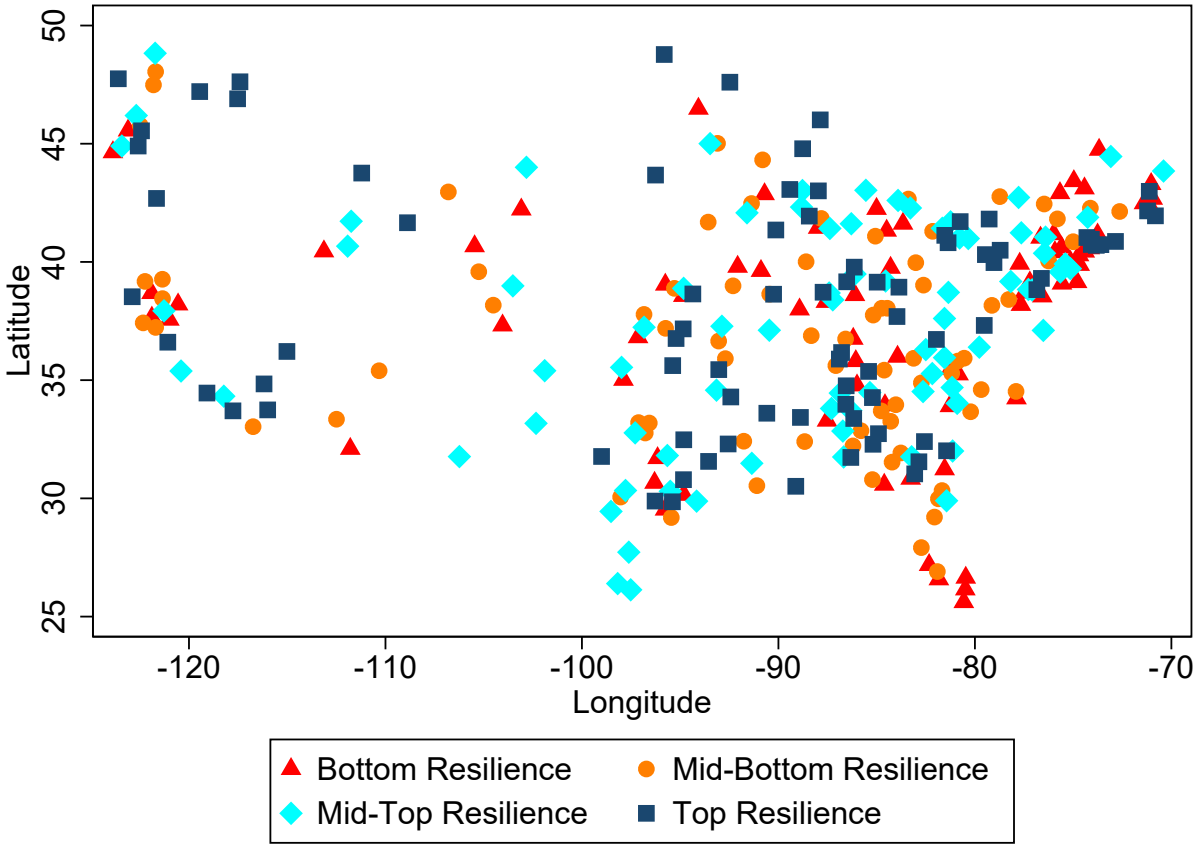
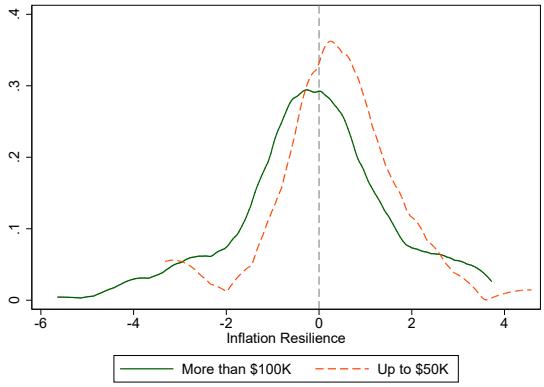


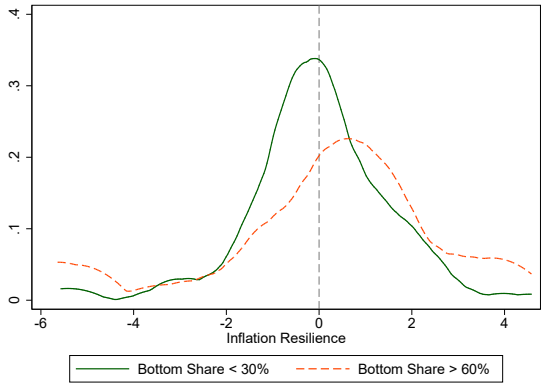
Figure 3: Distribution of Inflation Resilience by US Counties' Characteristics

Each panel of this figure compares the distribution of US county-level inflation resilience estimates across groups of counties by characteristics of the demand and the IO of local consumption that relate to the scope for trading down. Panel A compares high- and low-average-income counties. Panel B compares counties with a high or low share of bottom quality consumption at the beginning of our sample, i.e. prior to the 2021-2022 inflation surge. Panel C compares counties with a high or low average number of brands available locally across NAICS-2 industries. Panel D compares counties with a high or low concentration of consumption transactions across local POIs.

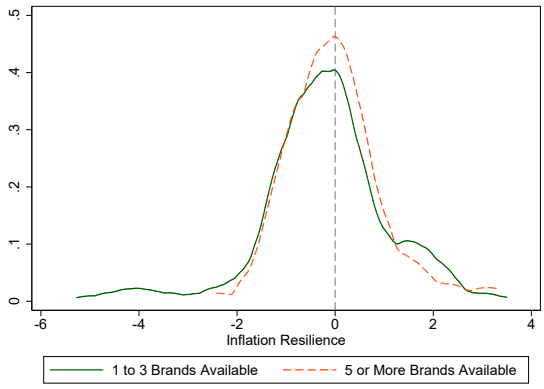
Panel A. Income



Panel B. Bottom Quality Share



Panel C. Number of Brands



Panel D. Concentration Transactions by POI

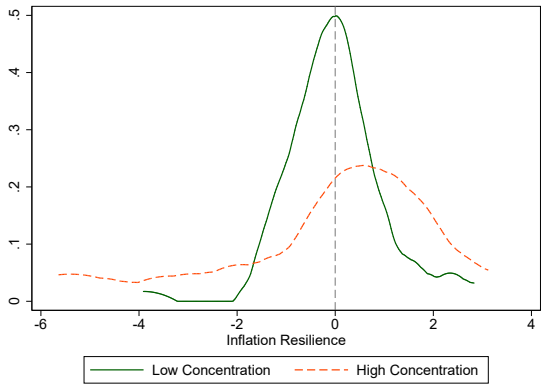
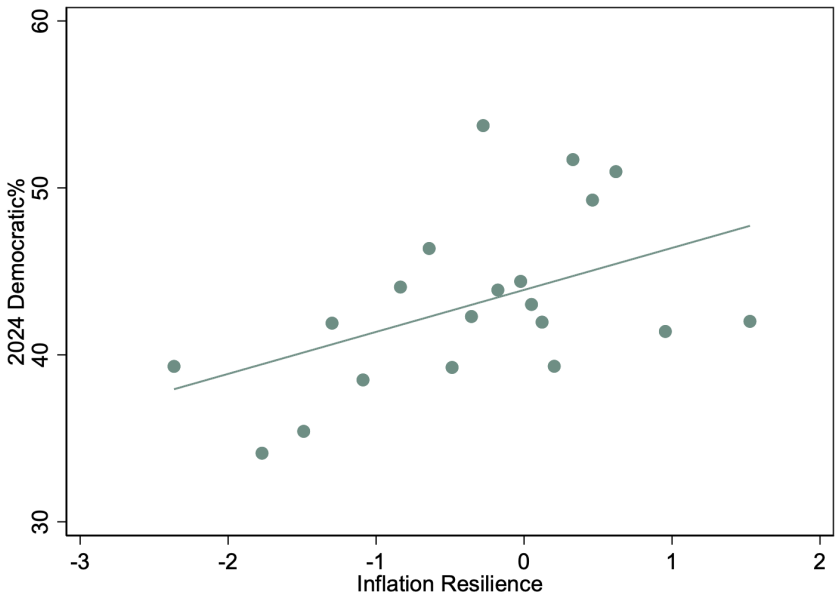


Figure 4: Inflation Resilience and Voting Behavior

This figure provides binned scatter plots of US county's voting outcomes against the standardized county-level estimates of inflation resilience. In Panel A, the y-axis measures the vote share of the Democratic Party candidate in the 2024 Presidential Election (Kamala Harris). In Panel B, the y-axis measures the difference between the vote share of the Democratic Party candidates in Presidential Elections in 2024 (Kamala Harris) and 2020 (Joe Biden). In both plots, state fixed effects are absorbed. Inflation resilience is trimmed at 1% - 99% level.

Panel A: Democratic Candidate Share (2024 Presidential Election



Panel B: Change in Democratic Candidate Shares (2024 - 2020 Presidential Elections)

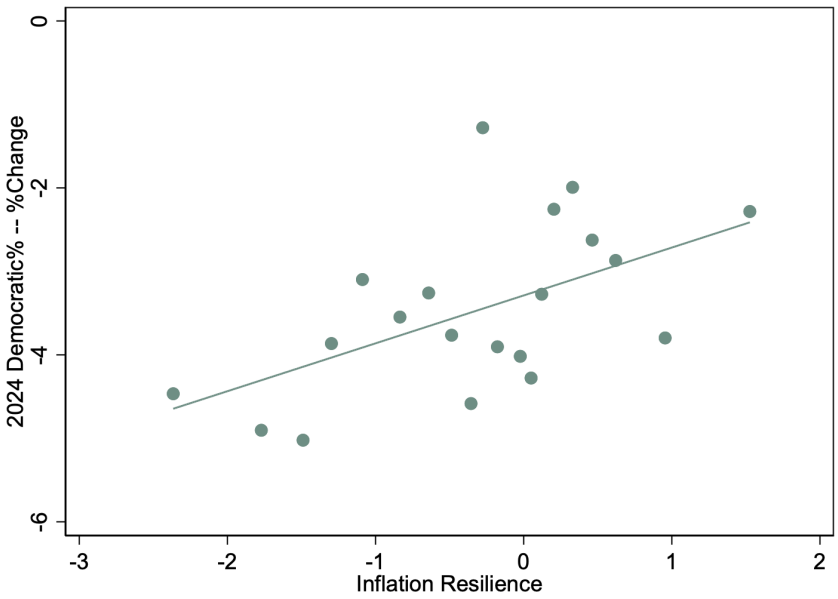


Table 1: Summary Statistics

This table summarizes our sample at the level of observation of our analysis (county-NAICS3-month). Average Spend represents the mean amount spent per transaction across the whole sample. N. Transaction is the total number of transactions. N. Consumers is the total number of consumers. Average Spend Low Quality, Average Spend Medium Quality, and Average Spend High Quality are the mean amounts spent per transaction for low-quality brands, medium-quality brands, and high-quality brands, respectively. PCE National Inflation is the national PCE inflation. CPI National Inflation is the national CPI inflation. Local Inflation is the regional CPI inflation. Our sample period runs from January 2020 to December 2022. All variables are winsorized at 1% level.

	Mean	SD	P10	P25	Median	P75	P90
Consumption Measures							
Average Spend	76.39	130.73	12.08	19.08	35.92	69.20	160.39
Average Spend Low Quality	42.18	53.70	9.88	13.94	25.29	47.08	84.45
Average Spend Medium Quality	73.25	109.72	13.53	19.45	39.30	70.61	160.82
Average Spend High Quality	138.45	243.80	21.57	31.77	57.88	119.39	313.54
N. Transaction	266.90	562.20	6.00	17.00	65.00	233.00	703.00
N. Consumers	154.32	283.94	5.00	14.00	48.00	155.00	413.00
Inflation Variables							
PCE National Inflation	3.42	1.73	1.16	1.55	4.01	5.19	5.37
CPI National Inflation	4.69	2.98	1.02	1.40	5.26	7.75	8.45
Local Inflation	4.68	3.10	0.99	1.40	5.30	7.40	8.70
N. Observations							44,686

Table 2: Inflation and Consumption Quality: Estimating Inflation Resilience

This table reports the estimates of the following linear specification by OLS:

$$Y_{c,j,t} = \beta \times \pi_{c,t} + \boldsymbol{X}_{c,j,t}\gamma + \epsilon_{c,j,t},$$

where for county c, industry j, and year-month t, $Y_{c,j,t} \in \{Bottom,\ Medium,\ Top,\ M-B,\ T-B\}$, denote the outcome variables indicated on top of each columns: $Bottom$, $Medium$, Top are respectively the consumption share of bottom-quality, medium-quality, and top-quality products; $M-B$ is the difference in the medium-quality and bottom-quality consumption shares; $T-B$ is the difference in the top-quality and bottom-quality consumption shares; $\pi_{c,t}$ (Local Inflation) is the regional CPI inflation. Demographics are from the ACS and include the local population, sex rate, median age, percent of African American residents, percent of White residents, percent of residents with a college degree, average family size, average household income, percent of homeowners with mortgages outstanding, percent of residents in the labor force, percent of residents 25-year-old or higher, median house value, median housing tax payments by residents.

	(1) Bottom	(2) Medium	(3) Top	(4) M-B	(5) T-B	(6) T-B	(7) T-B	(8) T-B
Local Inflation	0.14***	-0.06*	-0.07*	-0.20***	-0.21***	-0.18***	-0.68***	-0.50**
	(0.04)	(0.04)	(0.04)	(0.06)	(0.07)	(0.06)	(0.15)	(0.21)
R-squared	0.00	0.00	0.00	0.00	0.00	0.01	0.47	0.47
Demographics	No	No	No	No	No	Yes	Yes	Yes
NAICS3 FE	No	No	No	No	No	No	Yes	Yes
County FE	No	No	No	No	No	No	Yes	Yes
Year-month FE	No	No	No	No	No	No	Yes	Yes
State × N. Inf	No	No	No	No	No	No	No	Yes
Observations	44686	44686	44686	44686	44686	44346	44344	44344

Standard Errors Clustered at State × Year-Month Level in Parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Table 3: Inflation Resilience: Margins of Adjustment

This table reports the estimates of the following linear specification by OLS:

$$Y_{c,j,t} = \beta \times \pi_{c,t} + \boldsymbol{X}_{c,j,t}\gamma + \epsilon_{c,j,t},$$

where for county c, industry j, and year-month t, $Y_{c,j,t}$ is the log average cost per transaction in Panel A, the log number of transactions per consumer in Panel B, and the log number of consumers in Panel C. The top of each column indicates the group of observations for which these outcome variables are computed. *Bottom*, *Medium*, *Top* are respectively observations at the bottom, medium, and top of the consumption quality distribution; $M - B$ indicates variables computed for the medium group minus those computed for the bottom group; $T - B$ indicates variables computed for the top group minus those computed for the bottom group; $\pi_{c,t}$ (Local Inflation) is the regional CPI inflation. Demographics are from the ACS and include the local population, sex rate, median age, percent of African American residents, percent of White residents, percent of residents with a college degree, average family size, average household income, percent of homeowners with mortgages outstanding, percent of residents in the labor force, percent of residents 25-year-old or higher, median house value, median housing tax payments by residents.

	(1) Bottom	(2) Medium	(3) Top	(4) M - B	(5) T - B	(6) T - B	(7) T - B	(8) T - B
	Panel A: log(Cost per Transaction)							
Local Inflation	2.59***	2.05***	1.74***	-0.54***	-0.86***	-0.71***	0.59	0.93
	(0.28)	(0.34)	(0.39)	(0.12)	(0.19)	(0.15)	(0.50)	(0.64)
R-squared	0.01	0.00	0.00	0.00	0.00	0.03	0.38	0.38
	Panel B: log (Transactions per Consumer)							
Local Inflation	-0.18***	0.96**	0.45	0.04	0.11**	0.07*	-0.53***	-0.36*
	(0.06)	(0.49)	(0.42)	(0.04)	(0.04)	(0.04)	(0.14)	(0.18)
R-squared	0.00	0.00	0.00	0.00	0.00	0.01	0.45	0.45
	Panel C: log (N. Consumers)							
Local Inflation	-0.15	-1.11**	-0.60	-0.96***	-0.45**	-0.43**	-2.13***	-2.24***
	(0.50)	(0.50)	(0.44)	(0.19)	(0.23)	(0.21)	(0.62)	(0.76)
R-squared	0.00	0.00	0.00	0.00	0.00	0.01	0.45	0.45
Demographics	No	No	No	No	No	Yes	Yes	Yes
NAICS3 FE	No	No	No	No	No	No	Yes	Yes
County FE	No	No	No	No	No	No	Yes	Yes
Year-month FE	No	No	No	No	No	No	Yes	Yes
State × N. Inf	No	No	No	No	No	No	No	Yes
Observations	44686	44686	44686	44686	44686	44346	44344	44344

Standard Errors Clustered at State × Year-Month Level in Parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Table 4: Heterogeneity in Inflation Resilience: Demographics

This table reports the estimates of the following linear specification by OLS:

$$T - B_{c,j,t} = \beta \times \pi_{c,t} + \beta_H \times Top\ \ Group_c \times \pi_{c,t} + \beta_M \times Mid\ \ Group_c \times \pi_{c,t} + \boldsymbol{X}_{c,j,t}\gamma + \epsilon_{c,j,t},$$

where $T - B_{c,j,t}$ is the difference in the top-quality and bottom-quality consumption shares in county c, industry j and month t; $Top\ \ Group_c$ and $Mid\ \ Group_c$ are indicators for whether county c belongs to the top third or medium third of the distribution by the demographic characteristic indicated on top of each column, and zero otherwise. $\pi_{c,t}$ (Local Inflation) is the regional CPI inflation. Demographics are from the ACS and include the local population, sex rate, median age, percent of African American residents, percent of White residents, percent of residents with a college degree, average family size, average household income, percent of homeowners with mortgages outstanding, percent of residents in the labor force, percent of residents 25-year-old or higher, median house value, median housing tax payments by residents.

	(1)	(2)	(3)	(4)
	Related to Scope for Trading Down		Unrelated to Scope for Trading Down	
	Income	Bottom Share	Education	Age
Local Inflation	-0.47**	-0.07	-0.54**	-0.42*
	(0.22)	(0.16)	(0.22)	(0.23)
Local Inflation × Top Group	-0.22**	0.17*	-0.13	-0.16
	(0.09)	(0.09)	(0.13)	(0.14)
Local Inflation × Mid Group	0.02	0.20**	0.06	-0.10
	(0.07)	(0.08)	(0.07)	(0.11)
R-squared	0.47	0.76	0.47	0.47
Top, Bottom Group Dummies	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
NAICS3 FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes
Observations	44,264	44,264	44,264	44,264

Standard Errors Clustered at State × Year-Month Level in Parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Table 5: Heterogeneity in Inflation Resilience: IO Structure of Local Consumption

This table reports the estimates of the following linear specification by OLS:

$$T - B_{c,j,t} = \beta \times \pi_{c,t} + \beta_H \times Top\ Group_c \times \pi_{c,t} + \beta_M \times Mid\ \ Group_c \times \pi_{c,t} + \boldsymbol{X}_{c,j,t}\gamma + \epsilon_{c,j,t},$$

where $T - B_{c,j,t}$ is the difference in the top-quality and bottom-quality consumption shares in county c, industry j and month t; $Top\ \ Group_c$ and $Mid\ \ Group_c$ are indicators for whether county c belongs to the top third or medium third of the distribution by the feature of the local IO of consumption indicated on top of each column, and zero otherwise; $\pi_{c,t}$ (Local Inflation) is the regional CPI inflation. Demographics are from the ACS and include the local population, sex rate, median age, percent of African American residents, percent of White residents, percent of residents with a college degree, average family size, average household income, percent of homeowners with mortgages outstanding, percent of residents in the labor force, percent of residents 25-year-old or higher, median house value, median housing tax payments by residents.

	(1) Number Brands	(2) Number Brands (by NAICS2)	(3) HHI Transactions	(4) HHI Transactions (by NAICS2)
Local Inflation	-0.65***	-0.59***	-0.43*	-0.02
	(0.22)	(0.23)	(0.23)	(0.29)
Local Inflation × Top Group	0.24**	0.44***	-0.20*	-0.57***
	(0.08)	(0.17)	(0.11)	(0.20)
Local Inflation × Mid Group	0.04	0.03	-0.03	-0.45**
	(0.07)	(0.09)	(0.11)	(0.21)
R-squared	0.47	0.76	0.47	0.47
Top, Bottom Group Dummies	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
NAICS3 FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes
Observations	44,264	44,264	44,264	44,264

Standard Errors Clustered at State × Year-Month Level in Parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Table 6: Estimating Inflation Resilience: Instrumental-Variable Strategy

This table reports the results for estimating the following two-stage-least-squares (2SLS) specification:

$$\pi_{c,t} = b_1 \times Instrument_{i,t} + \boldsymbol{X}_{c,j,t}\gamma_1 + \epsilon_{1,c,j,t},$$
$$T - B_{c,j,t} = \beta_{iv} \times \widehat{\pi}_{c,t} + \boldsymbol{X}_{c,j,t}\gamma_2 + \epsilon_{2,c,j,t},$$

where $Instrument$ is obtained in the first stage as described in section 5.1; $Local\ Inflation$ is the regional CPI inflation. $Bottom$, $Medium$, and Top are respectively the consumption share of bottom-quality, medium-quality, and top-quality products. $T - B$ is the difference between the top-quality and bottom-quality consumption shares. $\pi_{c,t}$ (Local Inflation) is the regional CPI inflation. Demographics are from the ACS and include the local population, sex rate, median age, percent of African American residents, percent of White residents, percent of residents with a college degree, average family size, average household income, percent of homeowners with mortgages outstanding, percent of residents in the labor force, percent of residents 25-year-old or higher, median house value, median housing tax payments by residents.

	(1) Local Inflation	(2) Bottom	(3) Medium	(4) Top	(5) T - B
Instrument	0.28***				
	(0.03)				
Local Inflation		2.22***	-0.22	-2.00***	-4.22***
		(0.62)	(0.69)	(0.72)	(1.15)
First-stage F-stats			78.84		
Demographics	Yes	Yes	Yes	Yes	Yes
NAICS3 FE	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes	Yes
Observations	44344	44344	44344	44344	44344

Standard Errors Clustered at State × Year-Month Level in Parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Table 7: Robustness 1: National Quality Measures

This table reports estimates of the specifications in Table 2 and Table 6 after replacing the local measure of consumption quality with the national measures indicated on top of Panel A and Panel B. Local Inflation is the regional CPI inflation. Bottom, Medium, and Top are respectively the consumption share of bottom-quality, medium-quality, and top-quality products. T - B is the difference of the consumption shares in top-quality and bottom-quality products. IV is defined in equation (5). Demographics are from the ACS and include the local population, sex rate, median age, percent of African American residents, percent of White residents, percent of residents with a college degree, average family size, average household income, percent of homeowners with mortgages outstanding, percent of residents in the labor force, percent of residents 25-year-old or higher, median house value, median housing tax payments by residents.

	(1) Bottom	(2) Medium	(3) Top	(4) T - B	(5) T - B
	OLS				IV
	Panel A: Quality – Average National Cost				
Local Inflation	0.21***	-0.14**	-0.07	-0.28***	-0.26**
	(0.07)	(0.06)	(0.04)	(0.10)	(0.12)
Observations	157695	157695	157695	157695	157695
	Panel B: Quality – Brand Names Categorization by ChatGPT-4o				
Local Inflation	0.10*	-0.01	-0.09*	-0.19**	-0.31***
	(0.06)	(0.06)	(0.05)	(0.09)	(0.12)
Observations	134162	134162	134162	134162	134162
Demographics	Yes	Yes	Yes	Yes	Yes
NAICS3 FE	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes	Yes
State N. Inf	Yes	Yes	Yes	Yes	Yes

Standard Errors Clustered at State × Year-Month Level in Parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Table 8: Robustness 2: Subsamples and Alternative Measures

This table reports estimates of the specifications in Table 2 and Table 3 on the subsamples and with the variation indicated on top of Panel A and Panel B. Local Inflation is the regional CPI inflation. Bottom, Medium, and Top are respectively the consumption share of bottom-quality, medium-quality, and top-quality products. T - B is the difference of the consumption shares in top-quality and bottom-quality products. Demographics are from the ACS and include the local population, sex rate, median age, percent of African American residents, percent of White residents, percent of residents with a college degree, average family size, average household income, percent of homeowners with mortgages outstanding, percent of residents in the labor force, percent of residents 25-year-old or higher, median house value, median housing tax payments by residents. Coefficients in columns (2)-(4) are multiplied by 100 for easier interpretation.

	(1) T - B	(2) log(Cost per Trans.)	(3) log (Transactions)	(4) log (N. Consumers)
	Panel A: High T - B			
Local Inflation	-1.25***	-0.42	-0.32	-4.31**
	(0.48)	(1.30)	(0.36)	(1.76)
Observations	14618	14618	14618	14618
	Panel B: Quality Measure: Full Sample Average			
Local Inflation	-0.44**	0.08	-0.36*	-2.06***
	(0.21)	(0.63)	(0.18)	(0.75)
Observations	44163	44163	44163	44163
Demographics	Yes	Yes	Yes	Yes
NAICS3 FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes

Standard Errors Clustered at State × Year-Month Level in Parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$

Table 9: Beyond Consumption: Inflation Resilience and Voting Behavior

This table reports OLS estimates of the following linear specification:

$$\Delta Democratic\ Share_{2020\to2024,c} = \beta \times Inflation\ Resilience_c + \beta \times \bar{\pi}_c + \boldsymbol{X}_c\gamma + \epsilon_{c,j,t},$$

where $Democratic\ Share_{2020\to2024,c}$ is the difference between the Democratic candidate's vote share in 2024 and in 2020 Presidential Election in county c, $Inflation\ Resilience_c$ are the standardized county-level estimates of inflation resilience, $\bar{\pi}_c$ is the average of local inflation at the county level over our sample period, and X_c is the average of demographic from the ACS at the county level: local population, sex rate, median age, percent of African American residents, percent of White residents, percent of residents with a college degree, average family size, average household income, percent of homeowners with mortgages outstanding, percent of residents in the labor force, percent of residents 25-year-old or higher, median house value, median housing tax payments by residents.

	(1)	(2)	(3)
Inflation Resilience ($\hat{\beta}_c$)	0.41***	0.42***	0.28**
	(0.15)	(0.14)	(0.15)
Local Inflation		0.39	0.06
		(0.30)	(0.22)
R-squared	0.02	0.04	0.33
Demographics	No	No	Yes
Observations	301	301	301

Standard Errors Clustered at State Level in Parentheses
* $p < .10$, ** $p < .05$, *** $p < .01$