Anatomy of Countercyclical Price-Change Dispersion: HANK meets Non-homothetic Demand*

Minghao Li†
Peking University

October 2026

Abstract

This paper shows that the rise in price-change dispersion during the Great Recession was driven largely by diverging adjustments across quality tiers: prices of low-quality goods fell sharply, whereas those of high-quality goods remained stable. Consumers of low-quality goods became more price-sensitive, inducing sellers to cut markups. We develop a multi-sector, two-agent HANK model with menu costs and non-homothetic preferences in which cyclical variation in desired markups generates persistent dispersion in price changes. Quantitatively, the model matches the evidence and implies that countercyclical price elasticity produces procyclical markups, which dampen output responses to demand and monetary policy shocks.

Keywords: Cyclical dispersion, Non-homotheticity, HANK, TANK, Menu costs, Monetary policy
JEL Classification: E30, E31, E50

*I thank the editor (Francesco Lippi) and the anonymous referees for their constructive comments and guidance. I am grateful to Giuseppe Moscarini, Zhen Huo and Tony Smith for their invaluable guidance and support. For helpful comments and suggestions, I thank Olivier Coibion, Joachim Hubmer, Lerong Li, Xiangliang Li, Emi Nakamura, Michael Peters, Raphael Schoenle, Vincent Sterk, Kjetil Storesletten, Michael Weber, Daniel Xu, Donghai Zhang, and participants at various seminars and conferences. I also thank Alberto F. Cavallo for sharing his data with me when I worked on this project. I thank the Kilts Center for Marketing at the University of Chicago Booth School of Business for access to the NielsenIQ Retail Scan, NielsenIQ Homescan and Promo data. Researcher(s)' own analyses were calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analysing and preparing the results reported herein. Researcher(s)' own analyses calculated (or derived) based on data from National Promotion Reports, LLC, and marketing databases provided by the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business.

†Email: mhli@nsd.pku.edu.cn. China Center for Economic Research, National School of Development, Peking University, 100871 Beijing, China.

1. Introduction

Firms' pricing behaviour plays a central role in macroeconomic analysis. A well-known empirical regularity is the countercyclical dispersion of price changes — the distribution of price changes becomes more dispersed during recessions (Vavra, 2014). Existing explanations attribute this phenomenon either to exogenous fluctuations in volatility (Vavra, 2014) or to endogenously generated changes in volatility arising from learning or information frictions (Bachmann and Moscarini, 2012; Baley and Blanco, 2019).

This paper proposes a different explanation. We show that the rise in price-change dispersion during recessions is largely driven by systematic divergence in price adjustments across goods of different quality. During downturns, the prices of low-quality goods fall sharply, while the prices of high-quality goods remain stable. As a result, price adjustments diverge even within the same store and narrowly defined product category.

A simple example illustrates this pattern. During the Great Recession, Starbucks raised prices for some specialty beverages while simultaneously lowering prices for basic coffee drinks in order to retain more price-sensitive customers. Such pricing strategies reflect changes in consumers' price sensitivity during downturns and suggest that firms adjust their markups differently across products targeting different segments of demand.[1]

We document this mechanism using the NielsenIQ Retail Scanner (NRS) data, which provide near-universal coverage of products sold across a large number of stores together with detailed product characteristics. These data allow us to observe price changes for essentially all products within a store and product category, overcoming a key limitation of commonly used datasets such as the BLS microdata, where only a small number of products are sampled within each category and store.

Our empirical analysis reveals four novel facts. First, using a variance-decomposition approach, we show that the rise in price-change dispersion is overwhelmingly a within-store phenomenon. Approximately 85% of the overall increase in dispersion occurs across individual goods (UPCs) within the *same* store and narrow product category, rather than across different retailers and categories.

Second, this within-store-category dispersion is largely explained by a quality gradient. By proxying product quality with relative unit prices, we find that divergent price adjustments between high- and low-quality goods account for roughly 60% of the rise in within-store-category dispersion.

Third, inflation rates for high- and low-quality goods sharply diverge during economic downturns. Prices of low-quality goods fall substantially, while prices of high-quality goods remain stable. For example, a one-percentage-point increase in the local unemployment rate induces a 1.51 percentage-point decline in annualised monthly inflation of goods in the lowest quality group, compared to only a 0.44 percentage-point decline for the highest quality group. These patterns persist in the matched sample after controlling for observed wholesale cost changes.

[1] August 20, 2009, *The New York Times*: "Starbucks, in a Recession Risk, Raises some Prices."

Finally, we provide evidence that these divergent price adjustments are driven by changes in household price sensitivity. Households that predominantly purchase low-quality goods exhibit significantly greater price sensitivity during the recession. Relative to households that primarily purchase higher-quality goods, they spend less, obtain greater coupon savings relative to net purchase spending, buy more items on sale, and substitute towards store brands. These patterns indicate that sellers of low-quality goods face higher demand elasticities during downturns and respond by lowering their markups.

Guided by these empirical findings, we develop a multi-sector two-agent HANK model with menu costs and non-homothetic preferences. The key feature of the model is that both income elasticity and price elasticity vary with households' real income. This allows changes in aggregate demand to translate into systematic shifts in firms' desired markups across goods.

We emphasise a crucial conceptual distinction between desired markup *levels* and inflation *dynamics*. First, non-homothetic demand determines desired markup levels. Following a negative demand shock, real consumption declines, especially for hand-to-mouth households who disproportionately consume lower-quality goods. Because preferences are non-homothetic, lower consumption increases effective price elasticities. Firms selling lower-quality goods therefore optimally choose lower desired markups relative to firms producing higher-quality goods. This mechanism is entirely about markup levels and would arise even in a flexible-price environment.

Second, nominal rigidity maps these ideal markup levels into persistent *price changes*. Under staggered price adjustment, inflation is forward-looking and depends on the expected discounted path of desired markups. Consequently, a sustained decline in desired markups for low-quality goods generates gradual price reductions over time, producing protracted and heterogeneous deflation. In this way, non-homotheticity determines the cross-sectional differences in desired markup levels, while nominal rigidity governs how these level effects propagate into persistent dispersion in price changes.

We calibrate the model using NielsenIQ microdata to match steady-state expenditure shares and pricing moments. When subjected to aggregate demand shocks calibrated to the Great Recession, the model successfully replicates the non-targeted divergent inflation dynamics across quality tiers. Furthermore, the model explains roughly 88% of the observed increase in between-quality price-change dispersion. Non-homothetic preferences are central to these results: in an otherwise identical model with homothetic CES preferences, price elasticity remains constant and dispersion exhibits small cyclical variation. A separate, simplified flexible-price benchmark illustrates the distinction between markup levels and inflation dynamics. In log changes, inflation equals the contemporaneous change in ideal markups plus wage inflation. As ideal markups mean-revert, their growth rates become positive and are larger for low-quality goods. Consequently, inflation is higher for low-quality goods after the first period, opposite to the empirical ordering.

Finally, beyond matching the micro evidence, the model yields a novel macroeconomic implication: an endogenous stabilisation mechanism. During downturns, declining income increases households'

price sensitivity, inducing firms to reduce their markups. These procyclical markup adjustments act like a positive supply shock, shifting the aggregate supply curve downward and partially offsetting the contractionary effects of demand shocks. As a result, relative to a benchmark model with homothetic CES preferences, both real demand shocks and monetary policy shocks produce substantially smaller output responses.

This paper makes three main contributions. First, we document a new empirical fact: the countercyclical dispersion of price changes largely reflects divergence in price adjustments across goods of different quality within narrowly defined product categories. Second, we propose a mechanism linking non-homothetic demand and nominal rigidities that generates persistent inflation differences across goods. Third, embedding this mechanism in a heterogeneous-agent New Keynesian framework reveals a novel stabilisation channel through which countercyclical price elasticity dampens the macroeconomic effects of demand and monetary shocks.

Related Literature. This paper contributes to several strands of literature. First, it relates to work on time-varying microeconomic dispersion and volatility and their macroeconomic implications (Bloom (2009)). Within the literature on countercyclical price-change dispersion (Vavra (2014), Bachmann et al. (2019)), Vavra (2014) attributes the phenomenon to countercyclical shocks to idiosyncratic volatility, while Bachmann and Moscarini (2012) and Baley and Blanco (2019) generate it through endogenous learning dynamics. Using detailed NielsenIQ data, we provide a novel demand-side explanation based on heterogeneous price adjustments driven by countercyclical price elasticity.[2] Closely related, Munro (2021) documents countercyclical dispersion in store-level sales growth and proposes a representative-agent flexible-price model with countercyclical consumer search. In contrast, we study dispersion in price changes across products within stores and embed non-homothetic demand in a TANK framework to analyse monetary transmission.

Second, our paper contributes to the heterogeneous-agent New Keynesian (HANK) literature (e.g., McKay et al. (2016), Kaplan et al. (2018), Auclert (2019), Bilbiie (2020)). Standard HANK models largely treat firms' pricing decisions as independent of demand-side heterogeneity. By incorporating non-homothetic preferences into a two-agent New Keynesian (TANK) framework, our model generates heterogeneity in income and price elasticities that produces divergent inflation dynamics across quality groups and links supply-side pricing to demand heterogeneity.[3] We further show that endogenous markup adjustments dampen the aggregate response of output to both real and nominal demand shocks.[4]

Third, our paper relates to recent work studying the macroeconomic implications of non-homothetic preferences, particularly non-homothetic price elasticity.[5] Empirically, Argente and Lee (2021) show

[2]Our findings are consistent with Berger and Vavra (2019), who show that firms' endogenous responses generate time-varying price-change dispersion in exchange-rate pass-through.

[3]Lan et al. (2026) similarly emphasise interactions between households' MPCs and the price stickiness they face for monetary transmission and optimal policy design.

[4]Related insights appear in Olivi et al. (2024), though their analysis focuses on supply shocks under Calvo pricing.

[5]Matsuyama (2023) provides a survey of non-homothetic preferences in macroeconomics and trade. Many studies focus

that high-income households experienced larger declines in their cost of living during the Great Recession by trading down. We complement this evidence by focusing on the supply side, showing that firms producing low-quality goods reduce markups in response to increased price sensitivity. While macroeconomic analyses of price elasticity remain limited, Guimaraes and Sheedy (2011) study monetary non-neutrality assuming ex-ante heterogeneity in price sensitivity across households. More recently, Olivi et al. (2024) examine optimal monetary policy under supply shocks in a HANK model with non-homothetic income and price elasticities, while Mongey and Waugh (2025) endogenise non-homothetic price elasticity using discrete-choice demand with incomplete markets. Our paper differs by providing empirical evidence of time-varying price elasticity and by extending the generalised non-homothetic CES framework of Fally (2022) and Auer et al. (2024) to a dynamic New Keynesian model with menu costs.[6]

Finally, our paper builds on quantitative menu cost models (Golosov and Lucas (2007); Nakamura and Steinsson (2010); Midrigan (2011); Vavra (2014); Karadi and Reiff (2019); Alvarez et al. (2022, 2025)) by demonstrating that demand-driven markup cyclicality critically shapes aggregate stabilisation. We also complement empirical studies on heterogeneous pricing behaviour within industries or product categories (Gilchrist et al. (2017); Kim (2021); Jaravel (2019); Cavallo and Kryvtsov (2024), Chen et al. (2025),Sangani (2026a), Sangani (2026b)) by documenting systematic differences in cyclical inflation across narrowly defined quality tiers during the Great Recession.

Outline. The paper proceeds as follows. Section 2 describes the data and the construction of key variables. Section 3 presents the main empirical results. Section 4 presents the model. Section 5 studies firms' price setting with non-homothetic demand. Section 6 evaluates our model quantitatively. Section 7 explores the macroeconomic implications. Section 8 concludes.

2. Data and Measurement

This section describes the datasets used in our empirical analysis and explains how we measure price changes and product quality, two central concepts in our study. Additional details are provided in Appendix C.1.

2.1 Data

Our empirical analysis combines three datasets provided by NielsenIQ: price data from the NielsenIQ Scanner Dataset, detailed household shopping and spending data from the NielsenIQ Consumer Panel, and wholesale cost data from PromoData.

on non-homothetic income elasticity in growth (Boppart (2014), Matsuyama (2019), Comin et al. (2021)) or business cycles (Andreolli et al. (2024), Orchard (2025)). Non-homotheticity in product quality has been documented by Bils and Klenow (2001), Handbury (2021), Jaimovich et al. (2019), Faber and Fally (2022), and Ferraro and Valaitis (2025).

[6]Related evidence that price elasticity rises during downturns is provided by Stroebel and Vavra (2019).

NielsenIQ Retail Scanner Dataset. Our primary data source is the NielsenIQ Retail Scanner Data (NRS), made available through the Kilts Center at the University of Chicago. The NRS covers approximately 40,000 stores across 90 retail chains, spanning 371 metropolitan statistical areas (MSAs) and 2,500 U.S. counties since 2006. It provides weekly pricing and sales information for every Universal Product Code (UPC) with recorded transactions in a given store and week. The dataset includes a wide range of product categories, such as food, non-food grocery items, health and beauty aids, and general merchandise. Compared to other scanner datasets, such as the IRI Symphony dataset, the NRS also provides richer information on product characteristics, which allows us to construct product quality measures.

The data are hierarchically organised into 10 departments, 125 product groups, 1,075 product modules, and 3.8 million UPCs.[7] For example, a UPC for milk belongs to the product module "Dairy-Milk-Refrigerated," the product group "Milk," and the department "Dairy." Throughout this paper, we use the terms UPC, good, and product interchangeably, as well as product module and product category. Product modules are narrowly defined; for instance, "orange juice" and "apple juice" are distinct modules. The homogeneity of UPCs within a product module, combined with detailed product-level attributes (e.g., package size, unit), allows us to construct meaningful measures of product quality.

Given the dataset's scale, we restrict our baseline analysis to a manageable subset. Following Kaplan et al. (2019), we focus on the Minneapolis-St. Paul Metropolitan Statistical Area (MSA) from 2006 to 2015. We retain balanced product modules over this period, resulting in a baseline sample of 282 stores across 16 retail chains and 907 product modules, encompassing approximately 380,000 UPCs.[8] To mitigate the influence of outliers, we drop UPCs with prices below 1 cent.

The baseline sample from 2006 to 2015 allows us to focus on the Great Recession. In the Appendix, we extend the empirical analysis along both the time and space dimensions to assess robustness. In Appendix Section C.5, we extend the NielsenIQ sample through 2023, which allows us to examine whether the patterns we document also arise during the 2020 recession and the 2021–2022 inflation surge. In Appendix Section C.6, we replicate our baseline empirical exercises in the Atlanta–Sandy Springs MSA. This market is geographically distinct yet comparable in size, making it well suited for replicating the full empirical analysis.

The NielsenIQ Consumer Panel and PromoData. We complement the NRS data with the NielsenIQ Consumer Panel (NCP), which provides household-level shopping information. The NCP is a panel survey of approximately 50,000 U.S. households since 2004. Projection weights are applied to ensure the sample is representative of the U.S. population. The dataset includes detailed household-level information on income, demographics, and shopping behaviour. Specifically, households report the

[7]The ten departments are: Health and Beauty Aids, Dairy, Dry Grocery, Frozen Foods, Deli, Packaged Meat, Fresh Produce, Non-Food Grocery, Alcohol, and General Merchandise.

[8]We exclude the Alcohol and General Merchandise departments, as these products are often purchased in store formats not tracked in the NRS data. Our findings likely generalise to other markets, as DellaVigna and Gentzkow (2019) show, using the NRS dataset, that U.S. retail chains adopt nearly uniform pricing strategies across different markets.

date of shopping trips, stores visited, quantities and prices of UPCs purchased, whether items were on sale, and coupon usage. The store identifiers in the NCP match those in the NRS, enabling us to link the two datasets. This linkage allows us to measure the quality of UPCs purchased by households, as detailed in Section 2.2.

Finally, we incorporate PromoData to obtain wholesale cost information at the UPC level. This dataset collects daily wholesale prices from wholesalers in 93 markets from 2006 to 2012, covering over 34,000 UPCs.[9] By merging PromoData with the NRS, we can control for wholesale costs when analysing retail price behaviour.

2.2 Measuring Product Quality

In our analysis, product quality is defined at the UPC level within a product module in a specific store. We measure quality using the *unit price* of each UPC, such as the price per gallon of refrigerated milk. The relative unit price across UPCs within the same product module and store serves as a proxy for quality. Intuitively, if customers choose a more expensive product (in terms of unit price) among highly homogeneous UPCs in the same store, we interpret this product as being of higher quality.[10]

Specifically, we define our quality measure $Q_{i,s,c}$ as the standardised unit price of UPC i in category c sold at store s:

$$Q_{i,s,c} = \frac{UP_{i,s,c} - \overline{UP}_{s,c}}{\sigma^{UP}_{s,c}}, \tag{1}$$

where $UP_{i,s,c}$ is the average unit price of UPC i over its sample period, $\overline{UP}_{s,c}$ is the average unit price of all UPCs in category c at store s, and $\sigma^{UP}_{s,c}$ is the standard deviation of unit prices across UPCs in category c at store s. This standardisation ensures that quality measures are comparable across product categories and stores.

Discrete quality groups are defined separately from this continuous index. In the rank-based exercises, we rank UPCs by their average unit price within each store-category cell, divide each rank by the maximum rank in that cell, and apply unweighted quantile cutoffs to the normalised ranks pooled across store–UPC observations in the relevant market and reference period. The cutoffs are not estimated separately within each cell, so the groups need not contain equal numbers of products within each cell or in the final matched sample. The number of groups and the reference period are specified for each exercise below.

A potential concern is that using unit prices to measure quality may introduce bias when analysing pricing behaviour. For instance, low-quality goods might exhibit higher inflation simply because their prices start at lower levels and adjust upward over time. To address this concern, we take the following steps. First, we use the average unit price of a product over its sample period to mitigate the influence of temporary price fluctuations. Second, we compute the correlation between the rank of a UPC's

[9] Each market has one confidential wholesaler, which may differ across markets.

[10] This approach is widely used in the macroeconomics and trade literature. See, for example, Argente and Lee (2021), Faber and Fally (2022), and Cavallo and Kryvtsov (2024). Our key distinction is that we define quality at the store level.

unit price in a given year and its rank four years later within the same store and product module. The rank correlation is 0.96 for unit prices, indicating that ranks are highly stable over time.

2.3 Measuring Price Changes

It is well documented that the price computation methods in the NRS introduce measurement errors in price changes.[11] To mitigate this bias, we focus on regular price changes by applying a sales filter to posted prices, following Midrigan (2011).[12] Details on the construction of this sales filter and its role in reducing measurement errors are provided in Appendix Sections C.1.2 and C.1.3.

For the dispersion analysis, we calculate weekly log changes in sales-filtered regular prices for each UPC-store combination. We pool the retained non-zero changes within each month and compute dispersion from these weekly observations.

Moments of regular price changes for the entire sample align with findings from prior studies using store scanner data (e.g., Midrigan (2011), Cavallo (2018)). Price adjustment patterns vary systematically across goods of different quality. Low-quality goods exhibit more frequent adjustments, larger absolute sizes, higher standard deviations of price changes, and more dispersed distributions of non-zero price changes, suggesting that low-quality goods have greater price flexibility.[13] To the best of our knowledge, this is the first study to systematically document differences in micro-level pricing moments across goods of different quality. The detailed summary statistics of regular price changes are presented in Appendix Table C.1.

3. Four Facts on Countercyclical Dispersion of Price Changes

Leveraging our merged datasets, we document four new facts that shed light on the drivers of countercyclical dispersion of price changes. Before presenting these findings, we first confirm that the price-change dispersion increased during the Great Recession, consistent with prior studies (e.g., Vavra (2014)). Appendix Figure C.3 illustrates the interquartile range (IQR) of price changes alongside the unemployment rate in the Minneapolis-St. Paul market from 2007 to 2015. The dispersion of price changes exhibits a clear countercyclical pattern: it rises during the recession and is positively correlated with the unemployment rate.[14]

[11] See, for example, the discussion in Eichenbaum et al. (2014).

[12] Regular price changes are generally considered more relevant for assessing macroeconomic shocks. See Kehoe and Midrigan (2015) for further discussion.

[13] For example, Alvarez et al. (2016) show that, in a large class of models, monetary non-neutrality is positively related to kurtosis and negatively related to the frequency of price changes. Hong et al. (2023) also find that the frequency of price changes is robustly negatively related to monetary non-neutrality.

[14] Appendix C.1.4 provides more details on the construction of price-change dispersion. Furthermore, Appendix Figure C.9 shows that price-change dispersion is also negatively correlated with local GDP. Additionally, in Appendix C.1.8, we show that the aggregate frequency of regular price changes is weakly countercyclical and positively correlated with price-change dispersion, consistent with findings in Vavra (2014).

Fact 1: The dominance of within-store-category dispersion

To further understand the sources of price-change dispersion, we adopt a variance-decomposition approach. Specifically, we decompose the total variance of regular price changes into two components: the within-store-category dispersion, reflecting variation in price changes within the same store and product category, and the between-store-category dispersion, reflecting variation in price changes across different stores and product categories.

Let j index a UPC-store-week observation, with UPC i, store s, and week w, and let $m(w) = t$ assign that week to month t. Define $dp_{j,s,c,t} = \log P^R_{i,s,c,w} - \log P^R_{i,s,c,w-1}$, where P^R is the weekly regular price and c is the product category. We pool retained non-zero changes within each store-category-month, giving each observation equal weight within its cell. A UPC-store combination may appear in several weeks of the same month. The observation-level decomposition is:

$$dp_{j,s,c,t} = \overline{dp}_{s,c,t} + [dp_{j,s,c,t} - \overline{dp}_{s,c,t}],$$

where $\overline{dp}_{s,c,t}$ is the arithmetic mean of retained weekly log price changes within a given store-category-month. Within-cell variation includes differences across products and across weeks for repeated observations of the same product. Using this decomposition, the total variance of price changes can be written as:

$$\mathrm{Var}(dp \mid t) = \underbrace{\mathrm{Var}_{\omega_{s,c,t}}\left(\overline{dp}_{s,c,t} \mid t\right)}_{\text{Between-store-category}} + \underbrace{\sum_{s,c} \omega_{s,c,t}\, \mathrm{Var}(dp_{j,s,c,t} \mid s, c, t)}_{\text{Within-store-category}}, \tag{2}$$

where $\omega_{s,c,t}$ is the annual revenue weight of each store-category cell, normalised over the cells observed in month t.[15] The first term measures dispersion in average price changes across store-category cells, while the second term measures the weighted average of price-change dispersion within store-category cells.

During the Great Recession, the rise in price-change dispersion was primarily driven by within-store-category variation. Panel (a) of Figure 1 plots the three terms in equation (2) for each month from 2007 to 2015. During the Great Recession (shaded area), approximately 85% of the increase in the overall variance of price changes stems from the within-store-category component, while the remaining 15% is attributed to the between-store-category component.[16]

[15] Appendix Section C.1.5 provides more details on the decomposition.

[16] Our result is robust to alternative orders of decomposition. Whether we first decompose variance at the store level and then at the product category level, or vice versa, the conclusions remain unchanged.

Figure 1: Variance Decomposition

(a) Between- and within store-category variance

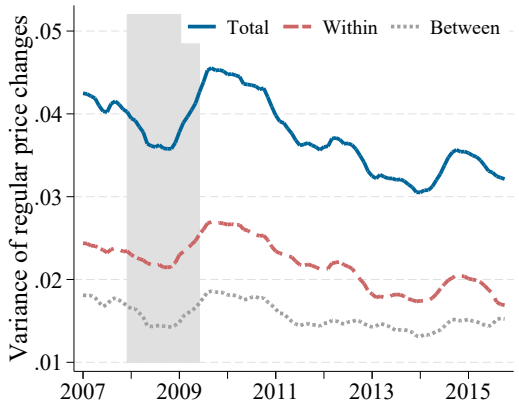

(b) Between- and within-quality-group variance

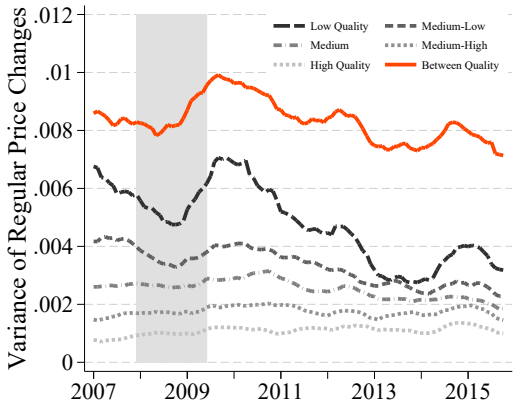

Notes: Each monthly statistic is calculated from retained weekly regular-price changes pooled within that month. A UPC-store combination may contribute multiple observations. Panel (a) shows decomposition of the total variance of price changes into within-store-product-category and between-store-product-category components based on equation (2). Panel (b) shows the decomposition of store-product-category variance into the between-quality components and within-quality components, using five groups formed by pooled quantile cutoffs of normalised within-store-category unit-price ranks over 2006–2015, as described in Section 2.2. The variance is aggregated using annual store-product-module weights. All plotted series are smoothed using a centred nine-month moving average (the current month and four months on each side). The shaded area denotes the Great Recession, December 2007–June 2009.

Fact 2: Between-quality dispersion explains most within-store-category dispersion

We now show that the increase in the within-store-category dispersion of price changes is primarily driven by diverging price adjustments between high- and low-quality UPCs. To illustrate this, we assign UPCs to five quality groups, ranging from low to high quality, using pooled quantile cutoffs of the normalised within-store-category unit-price ranks constructed over 2006–2015. These assignments are fixed over time; we do not impose equal group sizes separately within each store-category cell. We then decompose the total variance of price changes within each store-category cell into between-quality-group and within-quality-group components, where the between-quality-group dispersion reflects variation in price changes across different quality groups, and the within-quality-group dispersion reflects variation in price changes within each quality group.

Formally, we express this decomposition as:

$$\mathrm{Var}(dp \mid s,c,t) = \underbrace{\mathrm{Var}_{\omega_{Q|s,c,t}}\left(\overline{dp}_{Q,s,c,t} \mid s,c,t\right)}_{\text{Between-quality-group}} + \underbrace{\sum_{Q} \omega_{Q|s,c,t}\,\mathrm{Var}(dp_{j,s,c,t} \mid s,c,Q,t)}_{\text{Within-quality-group}}. \tag{3}$$

where $\omega_{Q|s,c,t} = N_{Q,s,c,t}/N_{s,c,t}$ is the share of retained weekly price-change observations in quality group Q within store s, category c, and month t. Here $N_{Q,s,c,t}$ and $N_{s,c,t}$ count weekly observations,

not distinct UPCs.[17] The term $\text{Var}_{\omega_{Q|s,c,t}}\left(\overline{dp}_{Q,s,c,t} \mid s,c,t\right)$ measures variance of group mean weekly price changes within store s and product category c in month t, while $\text{Var}(dp_{j,s,c,t} \mid s,c,Q,t)$ captures variance of retained weekly price-change observations within each quality group Q within store s and category c in month t. We aggregate this decomposition across store-category cells using the annual revenue weights.

Panel (b) of Figure 1 presents the decomposition results. The variance between quality groups rises sharply during the recession and accounts for approximately 62% of the overall increase in the within-store-category dispersion. The remaining 38% comes from dispersion within quality groups, with the largest contribution originating from the lowest-quality group.[18]

Fact 3: The inflation of high- and low-quality goods diverges during the recession

Fact 2 underscores the importance of understanding the pricing behaviour of high- and low-quality goods to explain the dynamics of price-change dispersion. We next show that the inflation rates of low-quality goods exhibit greater sensitivity to local unemployment rate fluctuations.

Empirical Approach. To analyse this, we use the five ordered quality groups in the precomputed store–UPC classification supplied with the regression inputs.[19] Group sizes are not rebalanced after the sample restrictions. We then estimate the following equation separately for each group.

$$\Delta p_{i,s,c,t} = \alpha_0 + \beta_0 u_t + \lambda_{f,s,c} + \epsilon_{i,s,c,t}, \tag{4}$$

where $\Delta p_{i,s,c,t} = 12 \times (p_{i,s,c,t} - p_{i,s,c,t-1})$ is the annualised monthly inflation rate of UPC i in store s and category c in month t, u_t is the local monthly unemployment rate, and $\lambda_{f,s,c}$ represents the firm-store-category fixed effects.[20]

The coefficient β_0 measures the correlation between inflation and the unemployment rate.[21] The magnitude of the estimated unemployment coefficient more than triples from the highest to the lowest quality group, as shown in Figure 2. Specifically, in the highest group, a one-percentage-point increase in the unemployment rate leads to a 0.44 percentage point decrease in the annualised monthly inflation rate, while in the lowest group, the same increase in unemployment results in a 1.51 percentage point decline in inflation.

A separate four-group exercise in Appendix C.4 provides evidence consistent with an extensive-

[17] Appendix Section C.1.5 provides more details on the decomposition.

[18] Since low-quality goods experience greater price reductions during the recession, the distribution of price changes becomes more symmetric compared to the distribution observed under positive trend inflation during normal periods. As a result, the price-change dispersion of low-quality goods increases during the recession.

[19] These assignments come from `quality_info.dta`. They are retained after the sample merges and are not replaced by the separately constructed five-group crosswalk used to match monthly price observations.

[20] The firm-store-category fixed effects are constructed by linking NielsenIQ Retail Scanner data with GS1 data, which assigns UPC codes to firms' products. This allows us to identify the firm associated with each UPC.

[21] Prior studies (e.g. Coibion et al. (2015)) have examined the cyclicality of aggregate posted prices but have not explored heterogeneity across quality tiers.

margin channel. In that exercise, low-quality goods experience a sharp decline in the frequency of positive price changes and a substantial rise in the frequency of negative price changes during recessions, while high-quality goods exhibit relatively stable pricing patterns. The figure compares the lowest and highest of four groups formed by pooled cuts of normalised within-store-category unit-price ranks. Because these groups differ from the five groups used above, this exercise provides supporting evidence rather than an exact decomposition of the regression coefficients.

Figure 2: Cyclicality of Inflation and Quality of Goods

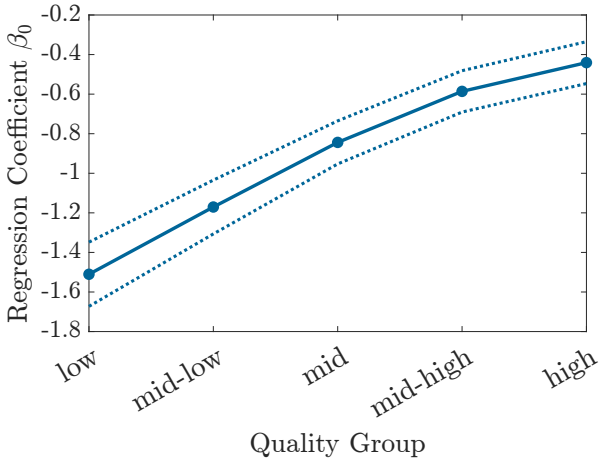

Notes: The five quality groups are the precomputed store–UPC assignments retained in the regression sample, as described in the text; their sizes need not be equal. This figure shows separate estimates of regression (4) in each group. The dependent variable is annualised monthly inflation, calculated from posted prices. All regressions include firm-store-category fixed effects, and standard errors are clustered at the store-category level. The dotted lines show the 95% confidence intervals.

Alternative Specification. To further validate our findings, we estimate an alternative specification that directly interacts the unemployment rate with a continuous quality index:

$$\Delta p_{i,s,c,t} = \alpha_1 + \beta_1 u_t \times Q_{i,s,c} + \gamma_1 Q_{i,s,c} + \gamma_2 u_t + \epsilon_{i,s,c,t}, \tag{5}$$

where $Q_{i,s,c}$ is the standardised quality index for UPC i within store s and category c, as defined in equation (1). The coefficient of interest, β_1, captures how inflation sensitivity to unemployment varies with product quality.

The results, presented in Table 1, confirm the quality-cyclicality relationship.[22] Column (1) shows that a one-percentage-point increase in unemployment lowers annualised monthly inflation by an additional 0.464 percentage points for a product whose quality is one standard deviation below the mean. This effect is sizeable, given that the average inflation reduction is 0.61 percentage points. Notably, annualised monthly inflation falls by 1.8 percentage points for goods at the 5th percentile

[22]The baseline regression uses sales weights, while Appendix Table E.1 demonstrates that our results are robust with quantity weights.

of quality, while it *rises* by 0.07 percentage points for those at the 95th percentile. This pattern aligns with anecdotal evidence, such as the "Starbucks story," where higher unemployment rates may even lead to price increases for high-quality goods. Column (2) includes store-category fixed effects to account for unobserved store-category heterogeneity. The estimates remain largely unchanged.

Table 1: Monthly Inflation, Unemployment Rate and Quality of Goods

	Monthly Inflation				
	(1)	(2)	(3)	(4)	(5)
Quality (Std)×	0.464***	0.455***	0.480***	1.69**	2.83***
Unemp Rate	(0.064)	(0.041)	(0.042)	(0.687)	(0.578)
Unemp Rate	-0.610***	-0.758***	-0.811***	-2.75***	-1.78***
	(0.058)	(0.034)	(0.035)	(0.751)	(0.391)
Quality (Std)×				-0.0006***	-0.0006***
Wholesale Inflation				(0.0001)	(0.0001)
Wholesale Inflation				-0.0046***	-0.0046***
				(0.0003)	(0.0003)
Fixed effects					
Store-Category		✓			
Firm-Store-Category			✓		✓

Notes: The dependent variable is the annualised monthly inflation. The quality index is standardised. Retail inflation and unemployment are seasonally adjusted; the wholesale cost-change measure in columns (4) and (5) is not. These columns also include an unemployment–wholesale cost-change interaction whose coefficient is not reported. Heteroskedasticity-robust standard errors are reported in columns (1) and (4); standard errors are clustered at the store-category level in columns (2), (3), and (5). All specifications include a constant that is not reported. The number of observations for columns (1) - (3) is 62,824,885. The number of observations for columns (4) and (5) is 1,185,046. *Significance*: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Controlling for Cost Factors. A potential explanation for the observed pattern is variation in wholesale costs. We address this in two ways. First, in Column (3) of Table 1, we introduce firm-store-category fixed effects, which allow us to compare price changes of UPCs produced by the same firm and sold in the same store and category. These fixed effects absorb persistent differences across firm-store-category units, but do not absorb time-varying firm-level cost shocks. The point estimates remain nearly identical to those estimated without firm fixed effects, showing robustness to these persistent differences. Second, Columns (4) and (5) control for UPC-level wholesale inflation, denoted $g_{i,t}$.

The full specification is

$$\pi^A_{i,s,c,t} = \alpha + b_{uQ} u_t Q_{i,s,c} + b_Q Q_{i,s,c} + b_u u_t + b_{gQ} g_{i,t} Q_{i,s,c} + b_g g_{i,t} + b_{ug} u_t g_{i,t} + \lambda_{f,s,c} + \epsilon_{i,s,c,t},$$

where $\pi^A_{i,s,c,t}$ is twelve times the seasonally adjusted monthly log retail price change, u_t is seasonally adjusted unemployment, and $Q_{i,s,c}$ is the standardised quality index. Column (4) sets $\lambda_{f,s,c} = 0$,

while Column (5) includes firm-store-category fixed effects.

In Table 1, Columns (4) and (5), the quality–unemployment interaction remains positive and statistically significant. These estimates establish robustness to the measured wholesale cost changes in the matched sample; they do not rule out unobserved cost shocks or separately identify demand-side effects.

Summary and Next Step. These findings clarify the nature of countercyclical price-change dispersion. During recessions, the inflation rates of high- and low-quality goods within the same store and category diverge, with low-quality goods experiencing substantial price declines while high-quality goods maintain stable prices. This pattern persists after controlling for observed wholesale cost changes in the matched sample, motivating an examination of demand-side factors while leaving open a role for unobserved cost shocks. This evidence sets the stage for further analysis of the mechanisms driving these demand-side effects.

Fact 4: Customers of low-quality goods become more price elastic during recessions

This section provides direct evidence that households primarily purchasing low-quality goods during the Great Recession become more price-sensitive. Compared to households that continue purchasing high-quality goods even during the downturn, these consumers reduce total spending, allocate a larger share of their expenditures to discounted goods, increase their purchases of private-label (generic) products, and obtain greater coupon savings relative to net purchase spending. These shifts in shopping behaviour suggest that sellers of low-quality goods face increasingly elastic demand during economic downturns, forcing them to lower markups as an optimal response.

Empirical Specification. To quantify households' purchasing behaviour, we first construct a quality index, $Q_h^{Recession}$, which measures the average quality of goods consumed by household h during the Great Recession:

$$Q_h^{Recession} = \sum_{t \in T_{Recession}} \omega_{h,t} Q_{h,t}, \tag{6}$$

where $\omega_{h,t}$ represents expenditure weights and $T_{Recession}$ includes quarters from Q4 2007 to Q2 2009. The index $Q_{h,t}$ measures the average quality of goods consumed by household h in quarter t (see Appendix C.1.6 for construction details). Households with lower $Q_h^{Recession}$ concentrated expenditures on low-quality goods during the recession. Notably, these households also tend to have lower and more cyclical incomes (see Appendix C.1.7 and Patterson (2023)).

For the regression analysis, we transform the recession quality index into a percentile rank across households. Let $\mathcal{H}_R$ denote the set of households in the household panel with a nonmissing recession quality index, and let $N_R = |\mathcal{H}_R|$. Define

$$R_h^{Recession} = \frac{\text{rank}_{j \in \mathcal{H}_R} \left(Q_h^{Recession} \right)}{N_R}, \tag{7}$$

where ranks are assigned in ascending order, with tied values receiving their average rank. Each household enters the ranking once. The ranking is unweighted; household projection weights are applied in the regressions.

The rank $R_h^{Recession}$ lies in $(0, 1]$, with higher values indicating that a household consumed higher-quality goods during the recession relative to other households. This rank is fixed over time for each household and is the quality measure used in the regressions below.

To examine how households' shopping behaviours respond to local unemployment fluctuations, we estimate

$$\begin{aligned} Y_{h,t} = \beta_0 + \beta_1 u_t + \beta_2 R_h^{Recession} \times u_t \\ + \gamma_1 X_{h,t} + \gamma_2 X_{h,t} \times u_t + \lambda_h + \epsilon_{h,t}. \end{aligned} \tag{8}$$

Here, $Y_{h,t}$ represents measures of household spending behaviour, u_t denotes the local unemployment rate, and $X_{h,t}$ contains time-varying household characteristics. Household fixed effects λ_h absorb time-invariant household characteristics, including the level of $R_h^{Recession}$. The coefficient of interest, β_2, captures the differential effect of unemployment on shopping behaviour between households that primarily purchase high- versus low-quality goods during the recession.

Table 2: Households' Recession Quality Index and Shopping Behaviour

	Dependent Variable							
	Log(Spending)		Share Deal		Share Generic		Share Coupon	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Unemployment Rate	-0.013***	0.045	0.017***	0.020	0.007***	0.011*	0.003***	0.007
	(0.005)	(0.063)	(0.001)	(0.017)	(0.001)	(0.006)	(0.000)	(0.006)
$R_h^{Recession}\times$ Unemployment Rate	0.028***	0.037***	-0.008***	-0.011***	-0.004***	-0.003***	-0.001	-0.002**
	(0.009)	(0.009)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
Household controls		✓		✓		✓		✓
Household controls× unemployment rate		✓		✓		✓		✓
Household FE	✓	✓	✓	✓	✓	✓	✓	✓
Observations	43339	43339	43339	43339	43339	43339	43339	43339

Notes: Standard errors are clustered at the household level. The regression is weighted by household projection weights provided by the NielsenIQ Consumer Panel. Even-numbered columns include household controls and interactions of all controls with the unemployment rate. Household controls comprise categorical indicators for the age and presence of children, the ages and education levels of the female and male household heads, marital status, and race, together with log household income adjusted for household size, which enters continuously. *Significance*: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Key Findings. Following Stroebel and Vavra (2019), we examine log total spending, expenditure shares on discounted and generic goods, and the ratio of coupon savings to net purchase spending as indicators of households' price sensitivity. Table 2 shows that, as unemployment rises, house-

holds with lower $R_h^{Recession}$ exhibit more negative spending responses and stronger shifts towards discounted and generic goods than households with higher ranks. These patterns are consistent with greater price sensitivity among consumers of lower-quality goods. Estimates for the coupon savings ratio point in the same direction but are less consistently significant. Alternative quality measures and specifications using recession and quality indicators (Appendix Tables E.2 and E.3) broadly preserve the direction of these differences, although statistical significance varies across outcomes and specifications.

Summary and Implications. Our findings are consistent with a demand-side contribution to countercyclical price-change dispersion. During downturns, households purchasing lower-quality goods exhibit shopping responses suggestive of a greater increase in price sensitivity. These patterns motivate a mechanism in which sellers of lower-quality goods face stronger incentives to reduce markups, contributing to divergent price adjustments across quality groups. The household evidence does not directly identify changes in demand elasticities or markups, but provides empirical motivation for the model developed below, which formalises this mechanism and assesses its quantitative implications.

4. A Business-Cycle Model with Non-homothetic Preferences

Motivated by our empirical analysis, we develop a multi-sector two-agent business cycle model that replicates key patterns observed in the data. On the demand side, to generate heterogeneous consumption baskets and time-varying price elasticity, we introduce a generalised non-homothetic CES preference that allows both the *income elasticity* and *price elasticity* of demand to vary with households' real income (Fally (2022),Auer et al. (2024)).[23]

On the supply side, to match the observed distribution of price changes, we incorporate a second-generation menu cost model featuring leptokurtic idiosyncratic productivity shocks and random menu costs. All aggregate shocks in the model are one-time and unanticipated.

4.1 Environment

Households. The economy is populated by a unit continuum of households, categorised into two types: hand-to-mouth households (H) and savers (S), with respective measures λ_H and λ_S.[24] Hand-to-mouth households lack access to financial markets and fully consume their income each period, while savers have full access to complete financial markets and own all firms in the economy. Each firm is managed by a risk-neutral agent.

A type-h household derives utility from real aggregate consumption C_t^h and labour supply N_t^h,

[23]The income elasticity of demand measures how quantity demanded responds to changes in real income, while the price elasticity of demand captures the sensitivity of quantity demanded to price changes.

[24]We adopt the Two-agent New Keynesian (TANK) framework over the Aiyagari-Bewley-Huggett-Imrohoroglu framework for tractability and computational feasibility, following Bilbiie (2020).

represented by the following utility function:

$$\mathcal{U}\left(\{C_t^h\}_{t=0}^{\infty}, \{N_t^h\}_{t=0}^{\infty}; \{\phi_t^h\}_{t=0}^{\infty}\right) = \sum_{t=0}^{\infty} \beta^t U\left(C_t^h, N_t^h; \phi_t^h\right), \tag{9}$$

where the flow utility is given by

$$U\left(C_t^h, N_t^h; \phi_t^h\right) = \frac{\phi_t^h \left(C_t^h\right)^{1-1/\vartheta}}{1-1/\vartheta} - \frac{\left(N_t^h\right)^{1+1/\gamma}}{1+1/\gamma}. \tag{10}$$

Here, ϑ denotes the elasticity of intertemporal substitution (EIS), and γ represents the Frisch elasticity. The demand shock ϕ_t^h, which shifts household h's marginal utility of consumption, follows an AR(1) process:

$$\phi_t^h = 1 - \rho_\phi + \rho_\phi \phi_{t-1}^h + \nu_t^h, \tag{11}$$

We prioritise demand shocks over supply shocks because, under nominal rigidities, supply shocks generate cost-push inflation and therefore imply counterfactual inflation dynamics during the Great Recession.

Firms. The production side of the economy consists of a single industry with a continuum of monopolistically competitive firms. To align with our empirical findings, we assume that goods are categorised into K quality groups (or sectors), ranked from low to high. Each firm ι in quality group k produces a differentiated good using labour $L_{k\iota t}$ via a linear production technology:

$$Y_{k\iota t} = Z_{k\iota t} L_{k\iota t}, \tag{12}$$

where $Z_{k\iota t}$ represents firm-specific productivity. To generate a fat-tailed distribution of price changes, we assume that idiosyncratic productivity evolves according to the following stochastic process:

$$\log Z_{k\iota t} = \begin{cases} \rho_k^z \log Z_{k\iota(t-1)} + \sigma_k^z \nu_{k\iota t}^z, & \text{with probability } \xi_k, \\ \log Z_{k\iota(t-1)}, & \text{with probability } 1-\xi_k, \end{cases} \tag{13}$$

where $\nu_{k\iota t}^z \sim N(0,1)$ is independent across firms.

Each period, a firm ι in quality group k can adjust its nominal price freely with probability ζ_k. Otherwise, it must pay a fixed menu cost $\overline{\kappa}_k$ (in units of labour) to make a price adjustment. The random menu cost $\tau_{k\iota t}$ follows:

$$\tau_{k\iota t} = \begin{cases} 0, & \text{with probability } \zeta_k, \\ \overline{\kappa}_k, & \text{with probability } 1-\zeta_k. \end{cases} \tag{14}$$

Monetary Policy. The central bank sets the nominal interest rate i_t according to a Taylor rule:[25]

$$i_t = (1+\bar{i})\left(\frac{1+\pi_t}{1+\pi^*}\right)^{\varphi_\pi}\left(\frac{Y_t}{Y^*}\right)^{\varphi_y}\left(\frac{V_t}{V^*}\right)^{\varphi_v}(1+u_t^i)-1, \tag{15}$$

where u_t^i represents shocks to the nominal interest rate and follows an AR(1) process:

$$u_t^i = \rho_u u_{t-1}^i + v_t^i. \tag{16}$$

Here, GDP is defined as $Y_t = (C_t^H)^{\lambda_H}(C_t^S)^{\lambda_S}$, and $\pi_t = P_t/P_{t-1} - 1$ is the aggregate inflation, where $P_t = (P_t^S)^{\lambda_S}(P_t^H)^{\lambda_H}$. The term V_t, introduced in equation (31), is a wedge in the Euler equation. Including V_t in the Taylor rule ensures tractability and determinacy. Period-t monetary shock is represented by v_t^i.

4.2 Non-homothetic CES Preferences

The aggregate real consumption of household h, denoted as C_t^h, is implicitly defined by a nested non-homothetic CES preference structure over all differentiated goods:

$$1 = \sum_{k=1}^{K}\left[g_k(C_t^h)\right]^{\frac{1}{\theta_t^h}}\left(C_{kt}^h\right)^{1-\frac{1}{\theta_t^h}}, \tag{17}$$

$$C_{kt}^h = \left(\int_{[k-1,k]}\left[C_{k\iota t}^h\right]^{\frac{\theta_t^h-1}{\theta_t^h}}d\iota\right)^{\frac{\theta_t^h}{\theta_t^h-1}}, \tag{18}$$

where the composite good C_{kt}^h is aggregated over differentiated goods $C_{k\iota t}^h$ within quality group k. The elasticity of substitution across varieties is given by θ_t^h, and the taste shifter for goods in quality group k is represented by $g_k(C_t^h)$, both depending on the real aggregate consumption C_t^h.

We assume that household h's price elasticity, θ_t^h, follows a log-linear relationship with real consumption:[26]

$$\log\theta_t^h = \delta_0 + \delta_1 \log C_t^h, \tag{19}$$

and specify the taste shifter $g_k(C_t^h)$ following Comin et al. (2021):

$$g_k(C_t^h) = \alpha_k (C_t^h)^{\varepsilon_k}. \tag{20}$$

Here, α_k is a demand shifter and ε_k affects the income elasticity of goods in quality group k, while δ_1 determines how price elasticity θ_t^h responds to changes in aggregate consumption. Since our empirical

[25]Traditional menu cost models often assume a money supply rule, $S = PY$. However, under demand shocks, this assumption produces counterfactual predictions, such as negative demand shocks leading to higher inflation.

[26]This specification ensures that θ_t^h is always positive. Nevertheless, our results remain qualitatively and quantitatively robust to alternative specifications, such as $\theta_t^h = \delta_0 + \delta_1 \log C_t^h$, as in Auer et al. (2024).

evidence suggests that price elasticity decreases as C_t^h rises, we assume that $\delta_1 < 0$. Notably, when $\theta_t^h = \theta$ is constant and $\varepsilon_k = 1-\theta$, this preference simplifies to the standard homothetic CES preference.

Lemma A.1 in Appendix A.1 provides sufficient conditions ensuring that the consumption aggregator C_t^h, defined over the consumption bundles $\boldsymbol{C}_t = \{C_{k\iota t} : k = 1, 2, \ldots, K, \iota \in [k-1, k)\}$, is well-behaved.

4.2.1 Non-homothetic Demand

Solving the household expenditure minimisation problem yields the Hicksian demand functions:[27]

$$C_{k\iota t}^h = \left(\frac{P_{k\iota t}}{P_{kt}^h}\right)^{-\theta_t^h} C_{kt}^h, \tag{21}$$

$$C_{kt}^h = \alpha_k (C_t^h)^{\varepsilon_k} \left(\frac{P_{kt}^h}{E_t^h}\right)^{-\theta_t^h}, \tag{22}$$

where E_t^h denotes household h's total nominal expenditure, $E_t^h = \sum_{k=1}^K P_{kt}^h C_{kt}^h$. The sectoral price index P_{kt}^h is given by:

$$P_{kt}^h = \left(\int P_{k\iota t}^{1-\theta_t^h} d\iota\right)^{\frac{1}{1-\theta_t^h}}, \tag{23}$$

which represents the minimal cost for household h to purchase one unit of C_{kt}^h. Under expenditure minimisation, the expenditure function for household h is given by:

$$\mathcal{E}^h(C_t^h; \boldsymbol{P}_t) = \left(\sum_{k=1}^K \alpha_k (C_t^h)^{\varepsilon_k} \left(P_{kt}^h\right)^{1-\theta_t^h}\right)^{\frac{1}{1-\theta_t^h}}, \tag{24}$$

which is the minimal expenditure required to achieve utility level C_t^h, and it equals E_t^h. Here, $\boldsymbol{P}_t \equiv \{P_{k\iota t}\}$ denotes the set of prices for all differentiated goods. Household h's indirect utility function, $\mathcal{C}^h(E_t^h; \boldsymbol{P}_t)$, is defined as the inverse of the expenditure function (24).

The following lemma ensures that households always exhaust their budgets:

Lemma 1. *Suppose the conditions of Lemma A.1 hold. Let all variety prices be strictly positive, and suppose that the sectoral price indices in (23) are finite. Then, over the interior $(1, \overline{C})$ of the feasible range of the consumption aggregator, the expenditure function $\mathcal{E}^h(C_t^h; \boldsymbol{P}_t)$ is strictly increasing in C_t^h.*

Household h's expenditure share on goods in quality group k is given by

$$m_{kt}^h = \frac{P_{kt}^h C_{kt}^h}{E_t^h} = \alpha_k (C_t^h)^{\varepsilon_k} \left(\frac{P_{kt}^h}{E_t^h}\right)^{1-\theta_t^h}, \tag{25}$$

[27]The derivation of the Hicksian demand is provided in Appendix A.2, while the Marshallian demand is characterised in Appendix A.3.

which sums to one across all quality groups, $\sum_{k=1}^{K} m_{kt}^h = 1$.[28] Finally, household h's price index is defined as

$$P_t^h = \left(\sum_{k=1}^{K} \alpha_k (C_t^h)^{\varepsilon_k + \theta_t^h - 1} (P_{kt}^h)^{1-\theta_t^h} \right)^{\frac{1}{1-\theta_t^h}}, \tag{26}$$

satisfying the identity $P_t^h C_t^h = E_t^h$.

4.2.2 Income and Price Elasticity

The Hicksian demand function, along with the first-order derivative of the indirect utility function, implies that the income elasticity of demand for good ι in quality group k is given by

$$\psi_{k\iota t}^h = \frac{\partial \log\left(P_{k\iota t} C_{k\iota t}^h\right)}{\partial \log E_t^h} = \theta_t^h + \frac{\varepsilon_k + \delta_1 \theta_t^h (e_t^h - p_{k\iota t})}{\overline{\varepsilon}_t^h + \delta_1 \theta_t^h (e_t^h - \overline{p}_t^h)} \left(1 - \theta_t^h\right), \tag{27}$$

where $e_t^h = \log(E_t^h)$, $p_{k\iota t} = \log(P_{k\iota t})$, and $\overline{\varepsilon}_t^h = \sum_{k=1}^{K} m_{kt}^h \varepsilon_k$ and $\overline{p}_t^h = \sum_{k=1}^{K} \int m_{k\iota t}^h p_{k\iota t} \, d\iota$ denote the expenditure-share-weighted averages of $\{\varepsilon_k\}$ and log prices, respectively. It follows directly that the weighted income elasticity sums to one: $\sum_{k=1}^{K} \int m_{k\iota t}^h \psi_{k\iota t}^h d\iota = 1$.

Equation (27) reveals that, all else equal, higher values of ε_k and $p_{k\iota t}$ result in higher income elasticity $\psi_{k\iota t}^h$. The effect of ε_k on income elasticity has been extensively studied in Comin et al. (2021), and our model reduces to theirs when $\delta_1 = 0$.

However, in contrast to their framework, even if ε_k is uniform across quality groups, a negative value of δ_1 implies that households allocate a larger share of their expenditure to cheaper goods when their total spending declines—such as during recessions. This occurs because lower income increases households' price sensitivity (higher θ_t^h), making lower-quality goods, which are relatively cheaper, more attractive. This channel is absent in models with constant elasticity of substitution.

In our model, household h's elasticity of substitution between two goods ι and ℓ is given by

$$\frac{\partial \log\left(C_{k\iota t}^h / C_{k\ell t}^h\right)}{\partial \log(P_{k\ell t} / P_{k\iota t})} = \theta_t^h = \exp\left(\delta_0 + \delta_1 \log C_t^h\right), \tag{28}$$

which depends on the household's real aggregate consumption.

Additional details on the properties of the non-homothetic CES preference are provided in Appendix A.

4.3 Household Problem

Optimisation Problems. Hand-to-mouth households, who lack access to financial markets, maximise their current-period utility by choosing consumption and labour supply, subject to their budget

[28] Note that m_{kt}^h corresponds to the summand in (17) when substituting equation (22) into (17).

constraint:

$$\max_{\{C_t^H, N_t^H\}} U(C_t^H, N_t^H; \phi_t^H)$$
$$\text{s.t. } E_t^H = o^H W_t N_t^H + T_t^H.$$

In contrast, savers maximise their lifetime utility by choosing consumption, labour supply, and bond holdings, subject to their per-period budget constraints:

$$\begin{aligned} &\max_{\{C_t^S, N_t^S, B_{t+1}\}} \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t^S, N_t^S; \phi_t^S) \\ &s.t.\ E_t^S + B_{t+1}/(1+i_t) \leq o^S W_t N_t^S + B_t + D_t/(\lambda_S) + T_t^S, \quad t = 0, 1, 2, \ldots \end{aligned} \tag{29}$$

where β is the discount factor, o^h represents household h's labour efficiency, W_t is the nominal wage, B_{t+1} denotes holdings of a nominal bond paying at $t+1$, i_t is the net nominal interest rate, D_t denotes total firm profits, and T_t^S and T_t^H are lump-sum transfers.

The government taxes firm profits at rate τ_π and redistributes the proceeds to hand-to-mouth households, such that $\lambda_H T_t^H = -\lambda_S T_t^S = \tau_\pi D_t$. This transfer scheme generates different cyclicality in households' income, playing a key role in shock transmission in HANK models (Bilbiie (2020)).

Since the hand-to-mouth households' problem is static, we focus on the savers' optimisation problem, which can be decomposed into two separate sub-problems: (i) The intertemporal allocation of expenditures $\{E_t^S\}$, and (ii) The intratemporal allocation of expenditures across differentiated goods, as characterised in Section 4.2. Proposition A.1 in Appendix A.4 characterises the savers' problem, including the two-step separation result and a sufficient condition for a unique solution.[29]

Savers' Intertemporal Choices. Savers' optimal intertemporal allocation is governed by the Euler equation:

$$\frac{\partial U(C_t^S, N_t^S; \phi_t^S)}{\partial C_t^S} = \mathrm{E}_t \left[\frac{\beta(1+i_t)}{1+\pi_{t+1}^S} \frac{\partial U(C_{t+1}^S, N_{t+1}^S; \phi_{t+1}^S)}{\partial C_{t+1}^S} V_t \right], \tag{30}$$

where $\pi_{t+1}^S = P_{t+1}^S / P_t^S - 1$, and

$$V_t = (1+\pi_{t+1}^S) \frac{\partial C_{t+1}^S / \partial E_{t+1}^S}{\partial C_t^S / \partial E_t^S} = \frac{1-\theta_{t+1}^S}{1-\theta_t^S} \times \frac{\overline{\varepsilon}_t^S + \delta_1 \theta_t^S (e_t^S - \overline{p}_t^S)}{\overline{\varepsilon}_{t+1}^S + \delta_1 \theta_{t+1}^S (e_{t+1}^S - \overline{p}_{t+1}^S)}. \tag{31}$$

The term $\partial C_t^S / \partial E_t^S$ captures how indirect utility C_t^S varies with total expenditure E_t^S. As total expenditure changes, savers reallocate their spending across quality groups and become more price elastic. This creates a non-homothetic wedge V_t between the current marginal utility of consumption and the future one in the Euler equation (30), compared to its homothetic CES counterpart.

[29] An earlier derivation of the separation result can be found in Blackorby et al. (1978), with a recent application in Comin et al. (2021).

Households' Intratemporal Choices. The hand-to-mouth households' optimisation problem is characterised solely by intratemporal conditions on consumption allocation and labour supply. The intratemporal condition for household h's labour supply is given by

$$\frac{\phi_t^h (C_t^h)^{-1/\vartheta}}{(N_t^h)^{1/\gamma}} = \frac{\partial \mathcal{E}^h(C_t^h; \boldsymbol{P}_t)/\partial C_t^h}{o^h W_t}. \tag{32}$$

Transmission from Demand Shocks to Consumption. The transmission from demand shocks ϕ_t to household consumption operates in two steps. First, a negative preference shock lowers desired real consumption. For hand-to-mouth households, the adjustment occurs contemporaneously through the intratemporal labour supply condition and budget constraint, directly reducing nominal expenditures; for savers, it occurs intertemporally through the Euler equation. Second, under non-homothetic preferences, lower consumption increases effective price elasticities and shifts expenditures towards lower-quality goods. Thus, demand shocks reduce real consumption, raise price sensitivity, and reallocate demand across quality tiers.

4.4 Firms' Price-setting Problem

Savers delegate the firm to a risk-neutral manager with discount factor β. This manager normalises profits using the nominal wage rate W_t. Let P_{-1} and P denote firm ι's price in the previous and current period. A firm's optimisation problem can be formulated recursively as:

$$V_{kt}(P_{-1}, Z, \tau) = \max_P \left\{ \frac{D_{k\iota t}}{W_t} + \mathrm{E}\left[\beta V_{k,t+1}(P, Z_{+1}, \tau_{+1})\right] \right\}, \tag{33}$$

where

$$D_{k\iota t} \equiv \sum_{h=H,S} \lambda_h \left(P_{k\iota t} - \frac{W_t}{Z_{k\iota t}} \right) C_{k\iota t}^h - \mathbb{1}_{\{P_{-1} \neq P\}} \tau_{k\iota t} W_t.$$

Here, $C_{k\iota t}^h$ represents the demand for good ι in quality group k by household h, as given by equation (21). The term $\tau_{k\iota t} W_t$ captures the cost of price adjustments, incurred when the firm changes its price.

Let $\varrho_{kt}(P_{-1}, Z, \tau)$ denote the policy function for the firm's price choice in group k at period t, and let $\Pr(Z_{+1}|Z)$ represent the transition probability of idiosyncratic productivity. The distribution over idiosyncratic states of firms in group k at period t, denoted by $\Gamma_{k,t}(P_{-1}, Z, \tau)$, evolves according to:

$$\begin{aligned} \Gamma_{k,t+1}(\mathcal{B}, Z_{+1}, \tau_{+1}) = &\left[\zeta_k \mathbb{1}_{\{\tau_{+1}=0\}} + (1-\zeta_k)\mathbb{1}_{\{\tau_{+1}=\overline{\kappa}_k\}}\right] \\ &\times \int_{\{(P_{-1},Z,\tau):\varrho_{kt}(P_{-1},Z,\tau)\in\mathcal{B}\}} \Pr_k(Z_{+1}|Z)\, d\Gamma_{kt}(P_{-1}, Z, \tau), \end{aligned} \tag{34}$$

for all Borel sets $\mathcal{B} \subseteq \mathbb{R}$. This equation describes how the distribution of firms' prices and states evolves over time, incorporating the distribution of the next menu-cost state.

Since the economy experiences one-time, unexpected aggregate shocks, both the policy functions

ϱ_{kt} and the distributions Γ_{kt} vary over time rather than remaining stationary.

4.5 Equilibrium

Given the laws of motion for exogenous shocks (13), (14), the path of demand shocks (11), and monetary policy shocks (16), an equilibrium of the economy consists of:

1. Households' demand for differentiated goods $\{C^h_{k\iota t}\}$, given by condition (21) of the households' decision problem, along with household aggregate consumption C^h_t and $\{C^h_{kt}\}_{k=1}^K$ that satisfy equations (17) and (18).
2. Firms' value functions $\left\{\{V_{k,t}\}_{k=1}^K\right\}_{t=0}^\infty$ and policy functions $\left\{\{\varrho_{kt}\}_{k=1}^K\right\}_{t=0}^\infty$ that solve firms' price-setting problem in equation (33), along with the distributions over firms' individual states $\left\{\{\Gamma_{k,t}(P_{-1}, Z, \tau)\}_{k=1}^K\right\}_{t=0}^\infty$ that evolve according to (34).
3. Prices, wages, interest rates, and inflation, $\{P^h_{kt}\}_{k=1}^K$, P^h_t, W_t, r_t, i_t, π^h_t, which satisfy equations (32), (23), (26), (30), and the Fisher equation $1 + i_t = (1 + r_t)(1 + \pi^S_{t+1})$.
4. The labour market and bond market clear.
5. A monetary policy rule that follows the specification in equation (15).

Appendix D.1 provides details on the computational method to solve for the transition path of the perfect foresight equilibrium in response to unexpected demand shocks and monetary policy shocks.

5. Price Setting Meets Non-homothetic Demand

Before proceeding to the quantitative analysis, we first examine firms' optimal pricing under non-homothetic demand. To make the analysis tractable, in this section, we introduce two simplifying assumptions relative to the full model in Section 4. First, we abstract from idiosyncratic productivity shocks, so that all price-adjusting firms within a given quality group k set the same price $P_{k\iota t} = P^\#_{kt}$ and produce the same output $C_{k\iota t} = C^\#_{kt}$ for all $\iota \in [k-1, k)$. Second, we assume menu costs are sufficiently large that firms only adjust prices when doing so is costless. This is equivalent to imposing a Calvo price-setting structure, where firms in group k adjust prices with probability ζ_k.

To account for the divergent inflation dynamics between quality groups observed in the data, we derive Phillips curves at the quality-group level. A central object in this derivation is the *ideal markup*, which we now define and characterise.

5.1 The Ideal Markup

The ideal markup, denoted by M^*_{kt} for quality group k in period t, is defined as a firm's optimal markup if it could adjust its price freely each period while all other firms face nominal rigidities.[30]

[30]This concept parallels that of "ideal" or "notional" prices in the New Keynesian literature (e.g., Woodford (2003)).

In what follows, we assume the aggregate consumption of savers exceeds that of hand-to-mouth households, i.e., $C_t^S > C_t^H$. Note that firm k's sales share to household type H, denoted ω_{kt}^H, is given by

$$\omega_{kt}^H = \frac{\lambda_H C_{kt}^H / C_{kt}^S}{\lambda_H C_{kt}^H / C_{kt}^S + \lambda_S}, \quad \omega_{kt}^S = 1 - \omega_{kt}^H, \tag{35}$$

$$\frac{C_{kt}^H}{C_{kt}^S} = \left(\frac{C_t^H}{C_t^S}\right)^{\varepsilon_k} \frac{(E_t^H)^{\theta_t^H}}{(E_t^S)^{\theta_t^S}} \left(W_t M_{kt}^*\right)^{\theta_t^S - \theta_t^H}, \tag{36}$$

where equation (36) follows from equation (22). The following proposition characterises the ideal markup M_{kt}^*.

Proposition 1. *Given nominal expenditures $\{E_t^h\}$, the nominal wage W_t, and all other firms' prices, the ideal markup M_{kt}^* satisfies*

$$M_{kt}^* = \frac{\omega_{kt}^H(M_{kt}^*)\,\theta_t^H + \omega_{kt}^S(M_{kt}^*)\,\theta_t^S}{\omega_{kt}^H(M_{kt}^*)\,(\theta_t^H - 1) + \omega_{kt}^S(M_{kt}^*)\,(\theta_t^S - 1)}, \tag{37}$$

where ω_{kt}^h depends on M_{kt}^* through (35) and (36). A sufficient condition for a unique M_{kt}^* is

$$\frac{\lambda_H}{\lambda_S} \frac{\left(\theta_t^H - \theta_t^S\right)^2}{\left(\theta_t^S - 1\right)^2} W_t^{\theta_t^S - \theta_t^H} \left(\frac{\theta_t^H}{\theta_t^H - 1}\right)^{\theta_t^S - \theta_t^H - 1} \left(\frac{C_t^H}{C_t^S}\right)^{\varepsilon_k} \frac{(E_t^H)^{\theta_t^H}}{(E_t^S)^{\theta_t^S}} < 1. \tag{38}$$

Proposition 1 highlights the joint determination of a firm's markup and its customer composition, as captured by the fixed-point condition (37). A firm selling a larger revenue share to more price-sensitive hand-to-mouth households tends to set a lower markup. That revenue share, ω_{kt}^H, is itself a function of the markup via (35) and (36), creating a feedback loop between optimal markups and customer composition.[31]

Corollary 1. *If price elasticity is homothetic (i.e., $\delta_1 = 0$), then $\theta_t^H = \theta_t^S = \theta$. In that case, the optimal markup M_{kt}^* is constant for all firms and equal to $\theta/(\theta - 1)$.*

This corollary establishes that markup heterogeneity in our model arises from non-homothetic price elasticity. Under homothetic preferences—or non-homothetic CES preferences with constant price elasticity (e.g., Comin et al. (2021))—price elasticity is uniform across households, implying identical ideal markups for all firms. Hence, markup dispersion in our setting is driven by variation in price elasticity itself, not merely by differences in customer composition.[32]

The following proposition links the income elasticity parameters associated with group k ($\{\varepsilon_k\}$) with the customer base and ideal markup.

[31] A related insight is shared in a recent paper by Mongey and Waugh (2025).

[32] Appendix B.6.2 compares our mechanism to models based on Kimball preferences, which also generate variable markups (e.g., Kimball (1995), Klenow and Willis (2016), Gopinath and Itskhoki (2010), Baqaee et al. (2024)).

Proposition 2. *Suppose the sufficient uniqueness condition in Proposition 1 holds and $C_t^H < C_t^S$. If $\varepsilon_k < \varepsilon_j$ and $\delta_1 < 0$, then $\omega_{kt}^H > \omega_{jt}^H$ and $M_{kt}^* < M_{jt}^*$.*

Proposition 2 implies that goods preferred by lower-income households (i.e., those goods with lower ε_k) have lower ideal markups, while goods preferred by higher-income households command higher ideal markups.[33] Furthermore, lower-quality goods generate a greater share of their revenue from lower-income households.

Appendix B.6 presents additional testable implications of the model, such as the prediction that low-income households allocate a larger fraction of their expenditures to low-quality goods. These implications align closely with empirical evidence.

5.2 The Quality-group Phillips Curves

The following proposition presents the log-linearised Phillips curves for each quality group. In what follows, log-deviations from steady-state values are denoted by lowercase variables, and steady-state values are denoted by the same letter with bars.

Proposition 3 (Quality-group Phillips Curves)**.** *Under the Calvo pricing assumptions stated above, suppose steady-state inflation is zero in every quality group, $\bar{\pi}_k = 0$ for all k.[34] To first order around this steady state, the inflation rate of quality group k, denoted π_{kt}, satisfies*

$$\pi_{kt} = \kappa_k(\mu_{kt}^* + w_{kt}) + \beta\,\mathbb{E}_t\pi_{k,t+1}, \tag{39}$$

where μ_{kt}^* is the ideal markup for firms in quality group k and $w_{kt} = w_t - p_{kt}$. The parameter ζ_k is the Calvo probability of price adjustment, and the slope coefficient $\kappa_k = \frac{\zeta_k[1-(1-\zeta_k)\beta]}{1-\zeta_k}$ increases with ζ_k.

Iterating (39) forward shows that current inflation in group k is proportional to the present discounted value of current and future "real ideal prices" $\mu_{kt}^* + w_{kt}$. Here, μ_{kt}^* captures deviations of the ideal markup from steady state, while w_{kt} reflects deviations in real marginal costs. Persistent low expected ideal markups and real marginal costs therefore induce firms to reduce prices continuously, leading to protracted deflation.

In a standard New Keynesian model, $\mu_{kt}^* = 0$ because desired markups are constant. With non-homothetic preferences, by contrast, μ_{kt}^* varies when households' price elasticity changes or when the composition of consumers shifts, as shown in equation (37).

Empirically, we find persistent inflation differences across quality groups during the Great Recession after controlling for wholesale cost changes. In our model, this implies that ideal markups for low-quality goods must be persistently lower than those for high-quality goods. We examine

[33] Goods with lower ε_k in our model correspond to low-quality goods in the data.

[34] The zero-inflation assumption applies to the analytical Phillips-curve derivation. The quantitative model retains the calibrated monthly steady-state inflation rate $\bar{\pi} = 0.002$; the relationship between its inflation dynamics and the approximate Phillips curves is examined in Appendix D.4.

these heterogeneous responses of ideal markups across quality groups in Section 5.3, and confirm the mechanism quantitatively in Section 6.

Discussion: Sticky Prices vs. Flexible Prices. The model's inflation dynamics crucially depend on the assumption regarding price stickiness. Under flexible prices and common marginal cost equal to the nominal wage, the inflation rate for sector k, expressed as a log price change, mirrors contemporaneous changes in the ideal price:

$$\pi_{kt}^{\text{flex}} = \mu_{kt}^* - \mu_{k,t-1}^* + \pi_t^w = \Delta\mu_{kt}^* + \pi_t^w,$$

Thus, flexible-price inflation is determined by the contemporaneous change in the ideal markup and nominal wage inflation, rather than by an expected future path of desired prices.

Under sticky prices, by contrast, inflation is completely forward-looking. Proposition 3 shows that current inflation π_{kt} depends on the current and expected future *deviations* of the ideal markup μ_{kt}^* and real marginal cost w_{kt} from steady state. Section 6.3 illustrates this distinction using a simplified flexible-price benchmark; the comparison does not hold all other model features fixed.

5.3 Heterogeneous Responses of Ideal Markups

Equation (39) demonstrates the key role of the ideal markup changes in determining group-specific inflation rates. We next examine how the ideal markup of group k, M_{kt}^*, responds to fluctuations in households' real expenditures. To isolate the effect of demand on ideal markups and present the mechanisms transparently, we hold constant the markups of other quality groups while varying households' real expenditures E_t^H/W_t and E_t^S/W_t.[35] In particular,

Proposition 4. *Taking log deviations in households' real expenditures $e_t^H - w_t$, $e_t^S - w_t$ as given, the ideal markup μ_{kt}^* in quality group k is given by*

$$\mu_{kt}^* = \eta_k^H \left(e_t^H - w_t\right) + \eta_k^S \left(e_t^S - w_t\right), \tag{40}$$

where η_k^h measures the elasticity of group k's ideal markup to real expenditure of household type $h \in \{H, S\}$, given by

$$\eta_k^h = \frac{\lambda_h \overline{C}_k^h}{\eta_k^{\mu^*}} \Bigg[\underbrace{\left[(1 - \overline{\theta}^h) + \frac{\overline{\theta}^h}{\overline{M}_k^*} \right] \overline{\psi}_k^h}_{\text{Compositional effect}} + \underbrace{\left(\frac{1}{\overline{M}_k^*} - 1 \right) \delta_1 \overline{\theta}^h \frac{\partial \log C^h}{\partial \log E^h}}_{\text{Markup effect}} \Bigg]. \tag{41}$$

The term $\eta_k^{\mu^*}$, defined by equation (B.17) in Appendix B.5, is strictly positive under the maintained ordering $1 < \overline{\theta}^S < \overline{\theta}^H < 2\overline{\theta}^S$. The income elasticity $\overline{\psi}_k^h$ follows from equation (27), and the slope of the indirect utility

[35] Equations (36) and (35) show that these are precisely the aggregate variables affecting a firm's ideal markup.

function $\partial \log C^h / \partial \log E^h$, defined by equation (B.15) in Appendix B.5, is positive.

Proposition 4 establishes that firms' adjustments of ideal markups to household expenditure changes can be decomposed into two distinct channels: the *compositional effect* and the *markup effect*.

Compositional Effect. The compositional effect captures how changes in households' real expenditures alter the composition of a firm's customer base and, consequently, its ideal markup, holding constant households' price elasticity. For a normal good ($\overline{\psi}_k^h > 0$), the compositional effect is positive for savers and negative for hand-to-mouth households. For example, a decline in E_t^H (holding E_t^S fixed) reduces hand-to-mouth households' spending on good k, decreasing their share in the firm's customer base. Since these households are more price-sensitive, this lowers the average price elasticity, allowing the firm to raise its markup. By contrast, for an inferior good ($\overline{\psi}_k^h < 0$), the compositional effect reverses: it is negative for savers and positive for hand-to-mouth households.

This effect explains Starbucks' decision to raise prices of high-end products like Frappuccino during the Great Recession. As low-income consumers are disproportionately impacted by the recession, many switched to cheaper alternatives like McDonald's black coffee. Demand for Frappuccinos became concentrated among higher-income, less price-sensitive customers, enabling Starbucks to raise its markup.

Markup Effect. The markup effect captures changes in firms' ideal markups driven by shifts in households' price elasticity, holding consumption baskets fixed. This effect is always positive: as households' expenditures decline, their price elasticity increases, forcing firms to reduce markups. Intuitively, when households have less income to spend, they become more sensitive to price changes, giving firms less pricing power. This effect explains Starbucks' decision to lower the price of black coffee: following a recessionary shock, low-income consumers—who primarily buy black coffee—become more price-sensitive, compelling Starbucks to reduce its markup on black coffee.

The next corollary, which is a direct implication of Corollary 1, establishes that non-homotheticity in price elasticity is essential for generating cyclical changes in ideal markups.

Corollary 2. *When $\delta_1 = 0$, firms' ideal markups do not respond to changes in household expenditures. That is, $\eta_k^H = 0$ and $\eta_k^S = 0$.*

Net Effect. The sign of η_k^h depends on the relative strengths of the compositional and markup effects. Under the maintained assumptions $\overline{\theta}^H > \overline{\theta}^S$, $\delta_1 < 0$, and $\eta_k^{\mu^*} > 0$, the markup effect is strictly positive for both household types. The sign of the compositional effect, however, depends on both the household type and the income elasticity of good k. Therefore, the overall response of the ideal markup to a shock is generally a quantitative matter. In the calibrated model, the markup effect dominates the opposing compositional effect.

6. Quantification

This section quantifies the model developed in previous sections. We begin by describing the calibration of the model's steady state. We then show that the calibrated model replicates key patterns in pricing and spending behaviour during the recession. Finally, we highlight the critical role of non-homothetic preferences in generating these results.

6.1 Calibration and Steady-state Fit

We calibrate the model's parameters so that its steady-state moments match the empirical targets. Specifically, we discipline the key parameters using NielsenIQ microdata on household expenditures and price changes. Table 3 lists all calibrated parameters, while Table D.1 shows the model's fit to the data.

Preference. The model is calibrated at a monthly frequency, with the discount factor set to $\beta = 0.94^{1/12}$. The share of hand-to-mouth households λ_H is fixed at 0.5.[36] Following conventional values in the macroeconomic literature, we set both the elasticity of intertemporal substitution ϑ and the Frisch elasticity γ to 1. Using detailed consumption data, Auer et al. (2024) estimate the elasticity of substitution for high-income and low-income households at 4.4 and 6.8, respectively. We apply these estimates to equation (19) to calibrate δ_0 and δ_1, which govern non-homotheticity in price elasticity, ensuring that the steady-state elasticity of substitution for savers and hand-to-mouth households matches 4.4 and 6.8, respectively.[37] Finally, the steady-state value of the demand shifter ϕ_h is equal to 1.

Technology and Pricing. The number of quality groups K is set to 5, from low- to high-quality, consistent with our empirical analysis. The steady-state monthly inflation rate $\overline{\pi}$ is calibrated to 0.002 to match the average inflation rate from 1990 to 2007. The persistence of idiosyncratic productivity shock ρ_k^z is fixed at 0.7, a value commonly used in the menu-cost literature (e.g., Nakamura and Steinsson (2010)).

Policy. The central bank responds to inflation and the non-homothetic wedge V_t, with the Taylor rule coefficient $\varphi_\pi = 1.4$ and $\varphi_v = -1$. We calibrate the profit tax rate τ_π, used for lump-sum income redistribution, to target the income cyclicality of hand-to-mouth households. This parameter plays a crucial role in determining the aggregate effects of monetary policy, particularly through the redistribution channel (Auclert (2019), Bilbiie (2020), Patterson (2023)). We therefore set τ_π to 0.18 to

[36]The hand-to-mouth households and savers correspond to low-income and high-income households in the data. Consistent with findings in Patterson (2023), households with low income in our sample exhibit both lower steady-state expenditure levels and more cyclical income patterns. These features align closely with the hand-to-mouth behaviour in our model, where such households have limited access to financial markets and face more volatile disposable income.

[37]Note that in Lemma 1 we restrict real consumption to $C \in [1, \overline{C}]$. We relax this restriction in our quantitative model.

align the income cyclicality of hand-to-mouth households with the estimate of 1.64 from Patterson (2023), based on U.S. Census LEHD data.[38]

Internally Calibrated Parameters. The remaining parameters are calibrated to match detailed statistics from the NielsenIQ Consumer Panel and Retail Scanner Data. These parameters are listed in Table 3 panel (d). Following the HANK literature, we classify households with top-quartile income as savers and bottom-quartile as hand-to-mouth. The labour efficiency parameters $\{o_h\}$ for savers and hand-to-mouth are set to 1.42 and 0.94, respectively, to match the steady-state ratio of nominal expenditures between the two groups. Recall that households' expenditure shares across quality groups, defined in equation (25), depend on the taste shifter parameters $\{\alpha_k\}$ and $\{\varepsilon_k\}$ in the utility function. We therefore calibrate $\{\alpha_k\}$ and $\{\varepsilon_k\}$ to jointly match the hand-to-mouth's and savers' steady-state consumption shares across quality groups.

Finally, pricing-related parameters are calibrated to replicate the micro pricing moments documented earlier. The empirical size and distribution targets use retained weekly regular-price changes pooled over the sample, while the frequency target is aggregated to months as described in Appendix Table C.1. These parameters govern the cross-sectional distribution of price changes. For each sector k, there are four such parameters: the menu cost $\overline{\kappa}_k$, the standard deviation of idiosyncratic productivity shocks σ_k^z, the probability of zero menu cost ζ_k, and the probability of idiosyncratic productivity shocks ξ_k. We calibrate these parameters to match the average frequency, absolute size, kurtosis of price changes, and the fraction of small price changes.[39] These internally calibrated parameters are broadly comparable to those in recent literature, as shown in Table 3 panel (d). The unconditional volatility of idiosyncratic shocks is 0.074 for the lowest-quality goods and 0.051 for the highest-quality goods.[40] The calibration implies that the total cost of price adjustment in the economy is roughly 0.9 percent of steady-state monthly revenues, slightly above the estimate of 0.7 percent by Levy et al. (1997) using supermarket data.

[38] Patterson (2023) estimates earnings sensitivity to GDP by deciles of the MPC and income distribution, finding that income elasticity increases with MPC and decreases with income, ranging from 0 to 2 (Figure A9 in Patterson (2023)). In TANK, up to first-order approximation, the income cyclicality of the hand-to-mouth is equal to $1 + (1 - \tau_\pi/\lambda_H)/\gamma$.

[39] Recall that small price changes are defined as those smaller than (1/2)mean($|dp|$).

[40] Comparable estimates in Midrigan (2011) and Mongey (2021), using scanner data, report average idiosyncratic shock volatility of 0.053 and 0.04. The calibration in Midrigan (2011) is weekly. Adjusted to monthly frequency, the probability of idiosyncratic shocks is 0.11, and the standard deviation of idiosyncratic shocks conditional on adjustment is 0.16.

Table 3: Calibrated Parameters

Parameter	Value	Description/Related to	Target/Source
(a) Preference			
β	$0.94^{1/12}$	discount factor	convention
ϑ	1	elasticity of intertemporal subs.	convention
γ	1	Frisch elasticity	convention
ϕ_h	1	ss. demand shifter	convention
$[\delta_0, \delta_1]$	[1.8, -2.3]	Non-homothetic price elasticity	diff in s.s. H/S price elasticity (Auer et al. (2024))
(b) Technology and Pricing			
$\overline{\pi}$	0.0020	ss. inflation rate	avg. inflation rate
ρ_k^z	0.7	persistence of idiosyncratic shock	convention
(c) Policy			
φ_π	1.4	Taylor rule coefficient	convention
φ_y	0	Taylor rule coefficient	convention
φ_v	-1	Taylor rule coefficient	determinacy
τ_π	0.18	profit tax rate	income cyclicality (Patterson (2023))
(d) Internally Calibrated parameters			
$[o^H, o^S]$	[0.94, 1.42]	labour efficiency	ratio of H/S spending
α_k	[0.21, 0.33, 0.22, 0.16, 0.09]	preference for good k	consumption shares
ε_k	[-5.00, -3.70, -3.25, -2.42, -1.60]	income-elasticity parameter	consumption shares
$\overline{\kappa}_k$	[0.04, 0.04, 0.04, 0.04, 0.04]	size of menu cost	pricing moments
ξ_k	[0.20, 0.18, 0.16, 0.16, 0.14]	Prob. of idiosyncratic prod. shock	pricing moments
ζ_k	[0.09, 0.09, 0.08, 0.08, 0.08]	Prob. of zero menu cost	pricing moments
σ_k^z	[0.165, 0.15, 0.145, 0.145, 0.135]	s.d. of idiosyncratic prod.	pricing moments

Notes: parameters at the quality-group level (e.g., income elasticity, size of menu cost) are listed from those producing low-quality goods to those producing high-quality goods.

Model Fit: Steady State. Table D.1 presents the model's steady-state fit for both targeted and non-targeted moments. As shown in panel (a), the model nearly perfectly replicates savers' and hand-to-mouth households' expenditure shares across quality groups. Panel (b) lists the targeted pricing moments, which the model also matches fairly well. The strong fit on both expenditures and pricing moments demonstrates that our model, featuring non-homothetic preferences and menu costs, effectively captures key dimensions of both households' spending behaviour and firms' pricing behaviour. Additionally, the baseline model performs well on non-targeted pricing moments, such as the standard deviation of price changes. Unsurprisingly, it also replicates the distribution of the absolute size of price changes for each quality group reasonably well, as shown in panel (c).

6.2 Model Fit: Business Cycles

With the calibrated steady state in place, we now examine the model's business-cycle properties. We demonstrate that the calibrated model successfully replicates key empirical patterns during recessions.

To this end, we assume that the economy is hit by unexpected contractionary demand shocks, ν_1^H and ν_1^S, in period 1. We set the persistence parameter of the demand-shifter process to 0.8 and calibrate their magnitudes to match the observed decline in nominal expenditures for both hand-to-mouth households and savers during the Great Recession, as measured in the NielsenIQ data.[41]

Consumption and Inflation Responses. Figure 3 plots the impulse responses of key economic variables following the demand shocks depicted in panel (a).[42] Panel (b) shows the evolution of savers' and hand-to-mouth households' demand elasticity, θ_t^h. As in our empirical results, hand-to-mouth households become more price-elastic during a recession, driven by the greater cyclicality of their consumption, as shown in panel (c). On impact, savers' real consumption drops about 1%, while that of hand-to-mouth households drops about 7%.

Note that we do not explicitly target inflation dynamics in our calibration. Nevertheless, our model closely matches these non-targeted data moments. Specifically, panel (d) plots the annualised monthly inflation for goods in the lowest-quality group (solid lines) and highest-quality group (dashed lines), with blue lines representing inflation generated by our baseline model and red lines representing the data counterparts. Consistent with the data, inflation for the lowest-quality goods falls more than that of the highest-quality goods. Relatedly, panel (e) plots the difference in annualised monthly inflation between the lowest- and highest-quality groups over 12 months following the shock.[43] Our baseline model successfully replicates the observed difference in inflation dynamics. Finally, panel (f) reproduces our empirical results from Figure 2 in Section 3, which depicts the cyclicality of inflation across the five quality groups.[44] Overall, the strong fit in inflation dynamics demonstrates our model's ability to capture key interactions among household income, price elasticity, and firms' pricing decisions.

Countercyclical Dispersion of Price Changes. Our empirical analysis (Figure 1) shows that the rise in price-change dispersion during the Great Recession was primarily driven by within-store-category dispersion, with between-quality-group dispersion accounting for approximately 60% of this variation. Here, we show that our model not only reproduces the countercyclical nature of price-change dispersion—where recessions triggered by demand shocks lead to a surge in dispersion—but, more importantly, also accurately captures the substantial contribution of between-quality-group dispersion.

Mirroring Figure 1, Figure 4 presents the evolution of total price-change dispersion in our model, decomposed into between-group and within-group components. The between-group dispersion increases by 0.0015, while the total within-group dispersion increases by 0.0020, with the largest

[41] We quantify the decline in nominal expenditures by comparing nominal expenditures in 2010 to those in 2007 (pre- and post-Great Recession).

[42] Appendix Figure E.1 presents IRFs of aggregate consumption and inflation, nominal expenditures and labour supply.

[43] The data in panel (d) cover the period from January 2009 to August 2010. The data counterpart in panel (e) spans from January 2009 to December 2009, with its midpoint aligning with the trough of the Great Recession.

[44] Following Hazell et al. (2022), we define changes in the unemployment rate as the opposite of the changes in total labour supply to compute the slope of the Phillips curve (or the cyclicality of inflation).

Figure 3: Impulse Responses to Demand Shocks

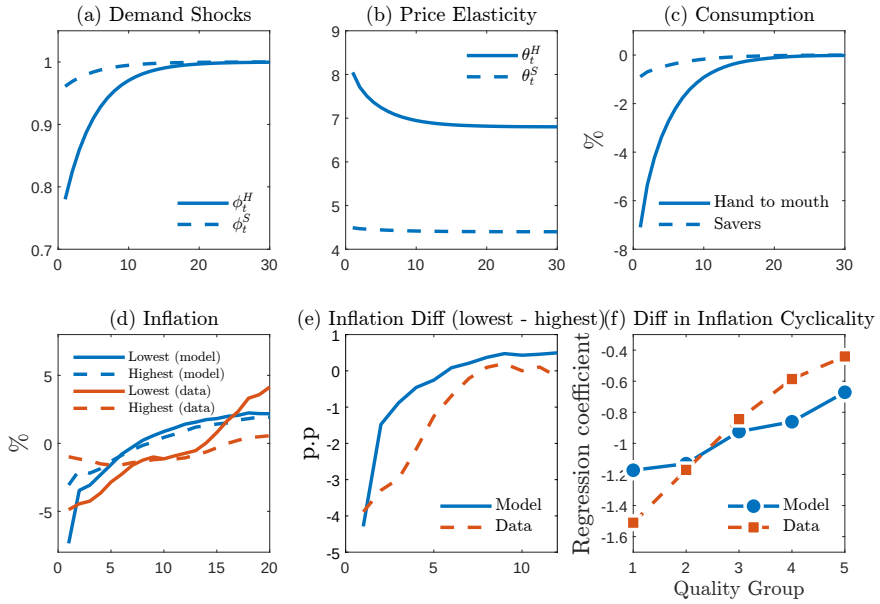

Notes: Horizontal axes denote months in panels (a) – (e). Shocks hit in month 1. Vertical axes in panel (c) denote percentage deviations from steady state. The data in panel (d) cover the period from January 2009 to August 2010. The data counterpart in panel (e) spans from January 2009 to December 2009, with its midpoint aligning with the trough of the Great Recession.

increase occurring among the lowest-quality goods. Thus, the between-group component accounts for roughly 43% of the overall increase in dispersion, compared to 60% observed in the data.[45] In terms of magnitude, our model explains 88% of the observed increase in the between-group variance. Specifically, in the data, the between-group variance increases by 0.0017 from November 2007 to its peak in September 2009, while in our model, this increase is 0.0015.

Figure 4: Countercyclical Dispersion of Price Changes

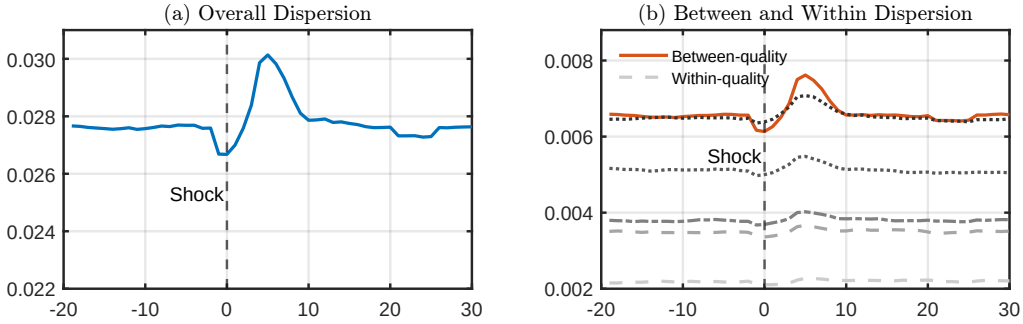

Notes: Horizontal axes denote months. Vertical axes denote price-change variance. The series are smoothed using a centred 5-month window. The overall dispersion, the between and within dispersion are averaged over 10 simulations.

[45]Guided by our empirical evidence, we do not model heterogeneity across retailers and product modules. As a result, the total variance of price changes in our model corresponds to the within-store variance.

6.3 The Roles of Non-homotheticity and Price Stickiness

We examine the role of non-homothetic preferences in accounting for the countercyclical dispersion of price changes and the divergent inflation dynamics across quality groups. We then use a simplified flexible-price benchmark to illustrate how nominal price adjustment affects inflation dynamics, subject to the additional changes in model features described below.

The Essential Role of Non-homotheticity in Price Elasticity. We show that the strong empirical fit of our model hinges on the assumption of non-homothetic preferences. To illustrate this, we construct an alternative version of the model in which we replace non-homothetic preferences with homothetic CES preferences ($\delta_1 = 0$; $\varepsilon_k = 1 - \exp(\delta_0)$), ensuring that all households share the same constant elasticity of substitution $\exp(\delta_0)$ and consume identical baskets of goods. This economy is subjected to the same demand shocks specified in Section 6.2.

The model results are presented in Appendix D.3.1. Unlike our baseline model, the homothetic version fails to generate enough countercyclical dispersion in price changes (Figure D.1). This is because homothetic preferences imply a constant price elasticity over time, and hence inflation differences and their cyclicality are too small compared to the data (Figure D.2). This leads to small changes in the between-group dispersion following the demand shocks.[46]

In Appendix D.3.2, we further show that our results remain largely unchanged when we eliminate non-homotheticity in income elasticity while retaining each household's baseline steady-state price elasticity and the baseline sensitivity of price elasticity to proportional consumption changes. This confirms that non-homotheticity in price elasticity is the critical mechanism driving our model's ability to replicate the empirical findings.

The Flexible-Price Benchmark. In Section 5.2, we discuss how inflation dynamics differ between sticky- and flexible-price environments. A simplified flexible-price benchmark illustrates how recovering markup levels can generate inflation patterns that differ from the empirical evidence.

Appendix D.4 presents this benchmark. It retains two household types, five quality groups, and non-homothetic demand, but removes idiosyncratic firm productivity and within-group price dispersion, normalises productivity to unity, and allows one price per quality group to adjust each period. The exercise therefore illustrates an inflation mechanism rather than providing a quantitative comparison that changes only nominal rigidity.

The two settings illustrate distinct inflation mechanisms. In the sticky-price model, inflation is forward-looking, as shown by the quality-group Phillips curve (39), and depends on current and expected future ideal markups and real marginal costs.[47] After a negative demand shock, the persistent and heterogeneous rise in households' price elasticity leads to a persistent, heterogeneous decline in ideal markups across quality groups, driving divergent deflation. In the simplified flexible-

[46] There are increases in between-group dispersion due to different degrees of nominal rigidity across quality groups.

[47] Figure D.5 in Appendix D.4 confirms that the quality-group Phillips curve holds approximately in our quantitative exercise.

price benchmark, log inflation equals contemporaneous log markup growth plus log wage growth. Markups recover after the initial shock, and the benchmark generates higher inflation for low-quality goods over the subsequent months shown in Figure D.4, contradicting the empirical evidence.

7. Aggregate Implications

In this section, we examine the aggregate implications of introducing non-homothetic price elasticity into a New Keynesian framework. Our main finding is that endogenous movements in price elasticity dampen the economy's aggregate output response to demand shocks. Relative to a benchmark model with homothetic CES preferences, both real demand shocks and monetary policy shocks generate smaller output fluctuations.

7.1 The Stabilising Effect of Cyclical Price Elasticity

A key feature of our model is that the price elasticity of demand is countercyclical: as household income declines, consumers become more price-sensitive. Firms respond by lowering their desired markups, generating procyclical markup movements.

Mechanism. In standard New Keynesian models, a decline in firms' markups lowers prices and acts as a positive supply shock. This shifts the aggregate supply (AS) curve downward, reducing inflation while increasing aggregate output. The stabilising mechanism in our model operates in a similar way, but with an important distinction: markup fluctuations arise endogenously from changes in households' price sensitivity rather than from exogenous markup shocks.[48]

Figure 5: Endogenous Shifts of the AS Curve

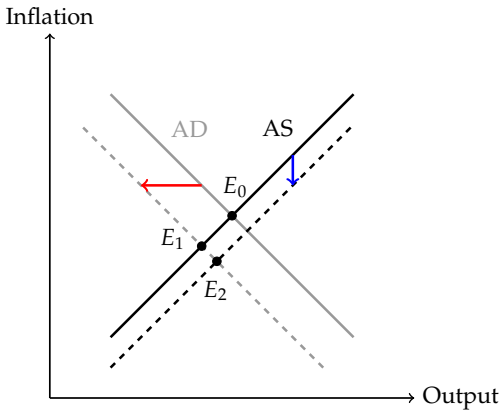

Figure 5 illustrates the mechanism using the classical AS–AD framework. The economy initially starts at equilibrium E_0. A negative demand shock—either real or nominal—shifts the AD curve

[48]For example, in Galí (2015), Chapter 5.2, markup shocks are introduced via exogenous time-varying elasticity of substitution.

leftward (from solid grey to dashed grey), generating a new equilibrium E_1 in a model with homothetic demand.

In our model, the same demand shock also increases households' price sensitivity, which lowers firms' desired markups and shifts the AS curve downward. The new equilibrium therefore moves to E_2, resulting in higher output and lower inflation relative to E_1. Conversely, during expansions, lower price sensitivity raises markups and shifts the AS curve upward, dampening the effects of positive demand shocks.

Overall, countercyclical price elasticity acts as an automatic stabiliser by attenuating the output response to aggregate demand fluctuations.

Quantification. We evaluate the quantitative importance of this stabilising mechanism by comparing our baseline model to an otherwise identical model with homothetic CES preferences outlined in Section 6.3, where income and price elasticities remain constant over time.[49]

Figure 6: Stabilising Effect: Demand and Monetary Shocks

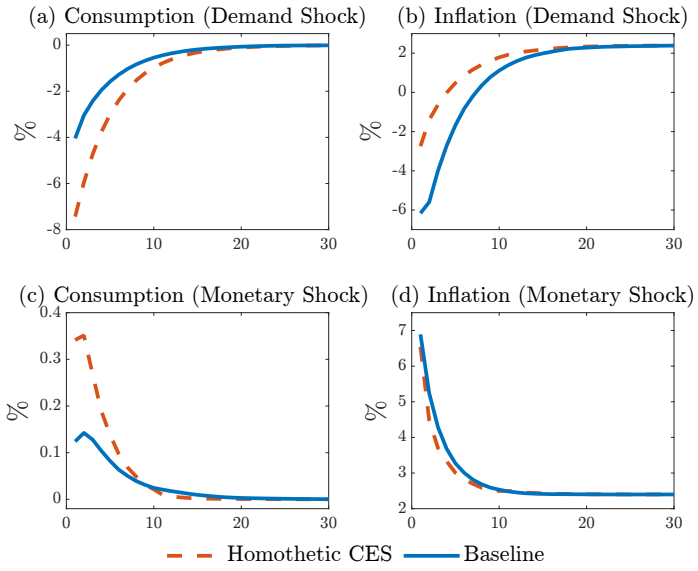

We first consider the demand shocks described in Section 6. Figure 6 panels (a) and (b) plot the impulse responses of aggregate consumption and inflation. Relative to the homothetic-CES benchmark, our model substantially dampens the output response: the cumulative response of aggregate consumption is -36.8% in the homothetic-CES model, while in our baseline it is -19.7%.[50] At the same time, inflation declines more in our baseline model because firms lower their markups in response to increased price sensitivity.

We next examine an expansionary monetary policy shock. Starting from each model's steady state, we set $\rho_u = 0.7$ in equation (16) and introduce an unanticipated innovation $\nu_1^i = -0.0025$ in

[49]Specifically, we set $\delta_1 = 0$ and $\varepsilon_k = 1 - \overline{\theta}$.

[50]Cumulative consumption responses sum monthly percentage deviations from steady state. Cumulative inflation responses sum annualised percentage-point deviations from steady-state inflation over the same 200-month horizon.

month 1, with $u_0^i = 0$ and no subsequent innovations. The same monetary shock path is applied to the baseline and homothetic-CES models. We compute the transition over 200 months and sum cumulative responses over this horizon. The cumulative consumption responses are 1.7% in the homothetic-CES model and 0.9% in the baseline model. The corresponding sums of annualised inflation deviations over 200 months are 10.5 and 13.2 percentage points, respectively. Figure 6 panels (c) and (d) display impulse responses of aggregate consumption and inflation for the first 30 months.

These stabilising effects are not driven by changes in the nominal rigidity channel—where variations in price elasticity affect the inaction region in menu-cost models—for two reasons. First, the baseline and homothetic-CES models generate similar steady-state frequencies and kurtoses of price changes. A generalisation of results in Alvarez et al. (2016) therefore implies similar real effects of monetary policy. Second, an expansionary monetary policy shock lowers price elasticity and would widen the inaction region, reducing the frequency of price changes and therefore implying larger output and smaller inflation responses. In contrast, our quantitative model predicts smaller output and larger inflation responses, indicating that the stabilising effect operates through endogenous markup adjustments rather than through changes in nominal rigidity.

Taken together, our model highlights a novel transmission mechanism: cyclical movements in price elasticity, and hence cyclical markups, create an endogenous stabilisation mechanism that dampens aggregate demand fluctuations.

8. Concluding Remarks

This paper examines the drivers of countercyclical dispersion in price changes, showing that the divergence in price adjustments across goods of different quality accounts for the majority of the increase in price-change dispersion during the Great Recession. Evidence from household shopping behaviour suggests that this phenomenon is driven by rising price elasticity for low-quality goods. Motivated by these empirical findings, we develop a multi-sector TANK model with non-homothetic preferences to explore its broader macroeconomic implications.

Our findings underscore the role of demand-side heterogeneity in explaining the heightened cross-sectional dispersion of economic variables during recessions, highlighting a promising avenue for future research. While the dataset used in this study offers a unique perspective on the research question, it primarily captures price changes in the retail sector and focuses on non-durable consumption goods. Extending these insights to other sectors will require further empirical investigation.

References

Acemoglu, D. (2008): *Introduction to modern economic growth*: Princeton University Press.

Algan, Y., O. Allais, and W. J. Den Haan (2008): "Solving heterogeneous-agent models with parameterized cross-sectional distributions," *Journal of Economic Dynamics and Control*, 32, 875–908.

Alvarez, F., A. Ferrara, E. Gautier, H. Le Bihan, and F. Lippi (2025): "Empirical investigation of a sufficient statistic for monetary shocks," *Review of Economic Studies*, 92, 2165–2196.

Alvarez, F., H. Le Bihan, and F. Lippi (2016): "The real effects of monetary shocks in sticky price models: a sufficient statistic approach," *American Economic Review*, 106, 2817–51.

Alvarez, F., F. Lippi, and A. Oskolkov (2022): "The macroeconomics of sticky prices with generalized hazard functions," *The Quarterly Journal of Economics*, 137, 989–1038.

Andreolli, M., N. Rickard, and P. Surico (2024): "Non-essential business-cycles," *Mimeographed, London Business School*.

Argente, D., and M. Lee (2021): "Cost of living inequality during the Great Recession," *Journal of the European Economic Association*, 19, 913–952.

Auclert, A. (2019): "Monetary policy and the redistribution channel," *American Economic Review*, 109, 2333–67.

Auer, R., A. Burstein, S. Lein, and J. Vogel (2024): "Unequal expenditure switching: Evidence from Switzerland," *Review of Economic Studies*, 91, 2572–2603.

Bachmann, R., B. Born, S. Elstner, and C. Grimme (2019): "Time-varying business volatility and the price setting of firms," *Journal of Monetary Economics*, 101, 82–99.

Bachmann, R., and G. Moscarini (2012): "Business Cycles and Endogenous Uncertainty," *Working paper*.

Baley, I., and A. Blanco (2019): "Firm uncertainty cycles and the propagation of nominal shocks," *American Economic Journal: Macroeconomics*, 11, 276–337.

Baqaee, D. R., E. Farhi, and K. Sangani (2024): "The supply-side effects of monetary policy," *Journal of Political Economy*, 132, 1065–1112.

Berger, D., and J. Vavra (2019): "Shocks versus Responsiveness: What Drives Time-Varying Dispersion?" *Journal of Political Economy*, 127, 2104–2142.

Bernanke, B., and O. Blanchard (2025): "What caused the US pandemic-era inflation?" *American Economic Journal: Macroeconomics*, 17, 1–35.

Bilbiie, F. O. (2020): "The New Keynesian cross," *Journal of Monetary Economics*, 114, 90–108.

Bils, M., and P. J. Klenow (2001): "Quantifying quality growth," *American Economic Review*, 91, 1006–1030.

Blackorby, C., D. Primont, and R. R. Russell (1978): *Duality, separability, and functional structure: Theory and economic applications*: North-Holland.

Bloom, N. (2009): "The impact of uncertainty shocks," *Econometrica*, 77, 623–685.

Boppart, T. (2014): "Structural change and the Kaldor facts in a growth model with relative price effects and non-Gorman preferences," *Econometrica*, 82, 2167–2196.

Cavallo, A. (2018): "Scraped data and sticky prices," *Review of Economics and Statistics*, 100, 105–119.

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

Chen, T., P. Levell, and M. O'Connell (2025): "Measuring cost of living inequality during an inflation surge," Technical report, Institute for Fiscal Studies.

Coibion, O., Y. Gorodnichenko, and G. H. Hong (2015): "The cyclicality of sales, regular and effective prices: Business cycle and policy implications," *American Economic Review*, 105, 993–1029.

Comin, D., D. Lashkari, and M. Mestieri (2021): "Structural change with long-run income and price effects," *Econometrica*, 89, 311–374.

DellaVigna, S., and M. Gentzkow (2019): "Uniform pricing in US retail chains," *The Quarterly Journal of Economics*, 134, 2011–2084.

Den Haan, W. J. (2010): "Comparison of solutions to the incomplete markets model with aggregate uncertainty," *Journal of Economic Dynamics and Control*, 34, 4–27.

Di Giovanni, J., Ş. Kalemli-Özcan, A. Silva, and M. A. Yldrm (2023): "Quantifying the inflationary impact of fiscal stimulus under supply constraints," in *AEA Papers and Proceedings* Volume 113, 76–80, American Economic Association.

Eichenbaum, M., N. Jaimovich, S. Rebelo, and J. Smith (2014): "How frequent are small price changes?" *American Economic Journal: Macroeconomics*, 6, 137–55.

Faber, B., and T. Fally (2022): "Firm heterogeneity in consumption baskets: Evidence from home and store scanner data," *The Review of Economic Studies*, 89, 1420–1459.

Fally, T. (2022): "Generalized separability and integrability: Consumer demand with a price aggregator," *Journal of Economic Theory*, 203, 105471.

Ferraro, D., and V. Valaitis (2025): "Consumption quality and employment across the wealth distribution," *Review of Economic Studies*, 92, 1801–1836.

Galí, J. (2015): *Monetary policy, inflation, and the business cycle: an introduction to the New Keynesian framework and its applications*: Princeton University Press.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek (2017): "Inflation dynamics during the financial crisis," *American Economic Review*, 107, 785–823.

Golosov, M., and R. E. Lucas (2007): "Menu costs and Phillips curves," *Journal of Political Economy*, 115, 171–199.

Gopinath, G., and O. Itskhoki (2010): "Frequency of price adjustment and pass-through," *The Quarterly Journal of Economics*, 125, 675–727.

Guimaraes, B., and K. D. Sheedy (2011): "Sales and monetary policy," *American Economic Review*, 101, 844–876.

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

Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson (2022): "The slope of the Phillips Curve: evidence from US states," *The Quarterly Journal of Economics*, 137, 1299–1344.

Hong, G. H., M. Klepacz, E. Pasten, and R. Schoenle (2023): "The real effects of monetary shocks: Evidence from micro pricing moments," *Journal of Monetary Economics*, 139, 1–20.

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

Jaravel, X. (2019): "The unequal gains from product innovations: Evidence from the US retail sector," *The Quarterly Journal of Economics*, 134, 715–783.

Kaplan, G., G. Menzio, L. Rudanko, and N. Trachter (2019): "Relative price dispersion: Evidence and theory," *American Economic Journal: Microeconomics*, 11, 68–124.

Kaplan, G., B. Moll, and G. L. Violante (2018): "Monetary policy according to HANK," *American Economic Review*, 108, 697–743.

Karadi, P., and A. Reiff (2019): "Menu costs, aggregate fluctuations, and large shocks," *American Economic Journal: Macroeconomics*, 11, 111–46.

Kehoe, P., and V. Midrigan (2015): "Prices are sticky after all," *Journal of Monetary Economics*, 75, 35–53.

Kim, R. (2021): "The effect of the credit crunch on output price dynamics: The corporate inventory and liquidity management channel," *The Quarterly Journal of Economics*, 136, 563–619.

Kimball, M. S. (1995): "The quantitative analytics of the basic neomonetarist model," *Journal of Money, Credit and Banking*, 27, 1241–1277.

Klenow, P. J., and O. Kryvtsov (2008): "State-dependent or time-dependent pricing: Does it matter for recent US inflation?" *The Quarterly Journal of Economics*, 123, 863–904.

Klenow, P. J., and J. L. Willis (2016): "Real rigidities and nominal price changes," *Economica*, 83, 443–472.

Knotek, E. S., and S. Terry (2008): "Alternative methods of solving state-dependent pricing models," *Federal Reserve Bank of Kansas City Working Paper*.

Lan, T., L. Li, and M. Li (2026): "Cyclical Inequality in the Cost of Living and Implications for Monetary Policy," *International Economic Review*, 67, 839–874.

Levy, D., M. Bergen, S. Dutta, and R. Venable (1997): "The magnitude of menu costs: direct evidence from large US supermarket chains," *The Quarterly Journal of Economics*, 112, 791–824.

Matsuyama, K. (2019): "Engel's law in the global economy: Demand-induced patterns of structural change, innovation, and trade," *Econometrica*, 87, 497–528.

——— (2023): "Non-CES aggregators: A guided tour," *Annual Review of Economics*, 15, 235–265.

McKay, A., E. Nakamura, and J. Steinsson (2016): "The power of forward guidance revisited," *American Economic Review*, 106, 3133–58.

Miao, J. (2020): *Economic dynamics in discrete time*: MIT Press.

Midrigan, V. (2011): "Menu costs, multiproduct firms, and aggregate fluctuations," *Econometrica*, 79, 1139–1180.

Mongey, S. (2021): "Market structure and monetary non-neutrality," Technical report, National Bureau of Economic Research.

Mongey, S., and M. E. Waugh (2025): "Pricing Inequality," Technical report, National Bureau of Economic Research.

Munro, D. R. (2021): "Consumer Behavior and Firm Volatility," *Journal of Money, Credit and Banking*, 53, 845–873.

Nakamura, E., and J. Steinsson (2010): "Monetary non-neutrality in a multisector menu cost model," *The Quarterly Journal of Economics*, 125, 961–1013.

Olivi, A., V. Sterk, and D. Xhani (2024): "Optimal Monetary Policy during a Cost of Living Crisis," Technical report, mimeo, UCL.

Orchard, J. (2025): "Non-homothetic Demand Shifts and Inflation Inequality," Finance and Economics Discussion Series 2025-085, Board of Governors of the Federal Reserve System.

Patterson, C. (2023): "The matching multiplier and the amplification of recessions," *American Economic Review*, 113, 982–1012.

Sangani, K. (2026a): "Cheapflation cycles," Technical report, National Bureau of Economic Research.

——— (2026b): "Complete pass-through in levels," *The Quarterly Journal of Economics*, 141, 1077–1135.

Stroebel, J., and J. Vavra (2019): "House prices, local demand, and retail prices," *Journal of Political Economy*, 127, 1391–1436.

Tauchen, G. (1986): "Finite state Markov-chain approximations to univariate and vector autoregressions," *Economics Letters*, 20, 177–181.

Vavra, J. (2014): "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *The Quarterly Journal of Economics*, 129, 215–258.

Woodford, M. (2003): *Interest and prices: Foundations of a theory of monetary policy*: Princeton University Press.

Young, E. R. (2010): "Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations," *Journal of Economic Dynamics and Control*, 34, 36–41.

Supplementary Appendix

Anatomy of Countercyclical Price-Change Dispersion:
HANK meets Non-homothetic Demand

Minghao Li
Peking University
mhli@nsd.pku.edu.cn

July 2026

A Non-homothetic CES preferences and Households' Optimisation **A-3**

A.1 Existence, Uniqueness, and Regularity of the Consumption Aggregator A-3

A.2 Deriving the Hicksian Demand A-7

A.3 Implicit Marshallian Demand A-8

A.4 Savers' Optimisation Problem A-8

B Proof of Theorems and Propositions **A-13**

B.1 Proof of Lemma 1 . . . . . A-14

B.2 Proof of Proposition 1 . . . . . A-16

B.3 Proof of Proposition 2 . . . . . A-18

B.4 Proof of Proposition 3 . . . . . A-20

B.5 Proof of Proposition 4 . . . . . A-21

B.6 Additional Results in Section 5 . . . . . A-23

C Appendix for Empirical Analysis **A-27**

C.1 Data A-27

C.2 Summary Statistics of Price Changes A-35

C.3 Countercyclical Dispersion of Price Changes A-35

C.4 Frequency and Size of Price Changes: Decomposition A-36

C.5 Inflation Dynamics during 2017–2023 . . . . . A-37

C.6 Empirical Results in Atlanta-Sandy Springs MSA A-39

C.7 Robustness A-41

D Appendix for Quantitative Analysis **A-43**

D.1 Numerical Solution Methods A-43

D.2 Steady-State Model Fit A-44

D.3 The Role of Non-homotheticity A-46

D.4 The Role of Price Stickiness A-49

E Additional Tables and Figures **A-54**

A. Non-homothetic CES preferences and Households' Optimisation

This section provides additional characterisation and properties of the non-homothetic CES preferences and households' optimisation problems. To simplify notation, we drop the household superscript h in this section.

A.1 Existence, Uniqueness, and Regularity of the Consumption Aggregator

To simplify notation, we suppress the household superscript h and the time subscript t. Let $\mathcal{I}_k = [k-1, k), \quad k = 1, \ldots, K$, and let $\boldsymbol{c} = \left\{c_{kj} > 0 : k = 1, \ldots, K,\ j \in \mathcal{I}_k\right\}$ denote a strictly positive consumption bundle. For a positive consumption bundle $\boldsymbol{c}$, define

$$\mathcal{H}(C; \boldsymbol{c}) \equiv \left[\sum_{k=1}^{K} \int_{\mathcal{I}_k} \left(\frac{c_{kj}}{G_k(C)}\right)^{\rho(C)} dj\right]^{\frac{1}{\rho(C)}}. \tag{A.1}$$

For a candidate aggregate-consumption level $C > 0$, define the following functions,

$$\begin{aligned} \theta(C) &\equiv \exp\left(\delta_0 + \delta_1 \log C\right), \\ \rho(C) &\equiv \frac{\theta(C) - 1}{\theta(C)}, \\ g_k(C) &\equiv \alpha_k C^{\varepsilon_k}, \\ G_k(C) &\equiv [g_k(C)]^{\frac{1}{1-\theta(C)}} = [\alpha_k C^{\varepsilon_k}]^{\frac{1}{1-\theta(C)}}. \end{aligned} \tag{A.2}$$

We establish the following lemma.

Lemma A.1. *Suppose that*

$$\alpha_k > 0, \qquad \sum_{k=1}^{K} \alpha_k = 1, \qquad \varepsilon_k < 0, \qquad \delta_1 < 0, \tag{A.3}$$

for every $k = 1, \ldots, K$. Restrict the admissible range of the consumption aggregator to $\mathcal{D} = [1, \overline{C}]$, where $\delta_0 + \delta_1 \log \overline{C} > 0$. It follows that $\theta(C) > 1$ and $\rho(C) \in (0, 1)$ for every $C \in \mathcal{D}$.

Define the natural admissible domain of strictly positive consumption bundles by

$$\mathfrak{B}_0 \equiv \left\{\boldsymbol{c} \gg 0 : \mathcal{H}(1; \boldsymbol{c}) > 1 > \mathcal{H}(\overline{C}; \boldsymbol{c})\right\}. \tag{A.4}$$

Then the following statements hold.

(i) For every $\boldsymbol{c} \in \mathfrak{B}_0$, there exists a unique $C(\boldsymbol{c}) \in (1, \overline{C})$ satisfying $\mathcal{H}(C(\boldsymbol{c}); \boldsymbol{c}) = 1$, or, equivalently,

$$1 = \sum_{k=1}^{K} \int_{\mathcal{I}_k} \left(\frac{c_{kj}}{G_k(C(\boldsymbol{c}))} \right)^{\rho(C(\boldsymbol{c}))} dj. \tag{A.5}$$

Thus, equation (A.5) defines a single-valued consumption aggregator $C : \mathfrak{B}_0 \longrightarrow (1, \overline{C})$.

(ii) The consumption aggregator is strictly increasing on $\mathfrak{B}_0$.

(iii) For every convex set $\mathfrak{B} \subseteq \mathfrak{B}_0$, the restriction of $C(\boldsymbol{c})$ to $\mathfrak{B}$ is quasi-concave.

Proof. The proof proceeds in five steps.

Step 1: Equivalence with the original nested aggregator. For any $C \in \mathcal{D}$, we have

$$G_k(C)^{-\rho(C)} = \left[g_k(C)^{\frac{1}{1-\theta(C)}} \right]^{-\frac{\theta(C)-1}{\theta(C)}} = g_k(C)^{-\frac{\theta(C)-1}{\theta(C)[1-\theta(C)]}} = g_k(C)^{1/\theta(C)}.$$

Therefore,

$$\sum_{k=1}^{K} \int_{\mathcal{I}_k} \left(\frac{c_{kj}}{G_k(C)} \right)^{\rho(C)} dj = \sum_{k=1}^{K} g_k(C)^{1/\theta(C)} \int_{\mathcal{I}_k} c_{kj}^{\rho(C)} dj.$$

Define the within-group composite

$$c_k(C) \equiv \left[\int_{\mathcal{I}_k} c_{kj}^{\rho(C)} dj \right]^{1/\rho(C)}.$$

Equation (A.5) is equivalent to

$$1 = \sum_{k=1}^{K} g_k(C)^{1/\theta(C)} c_k(C)^{1-1/\theta(C)},$$

which is exactly the nested consumption aggregator defined in equations (17) and (18).

Step 2: A normalised power-mean representation. Define

$$K_k(C) \equiv G_k(C) \alpha_k^{1/\rho(C)} = G_k(C) \alpha_k^{\frac{\theta(C)}{\theta(C)-1}}.$$

Using the definition of $G_k(C)$, we obtain

$$K_k(C) = (\alpha_k C^{\varepsilon_k})^{-\frac{1}{\theta(C)-1}} \alpha_k^{\frac{\theta(C)}{\theta(C)-1}} = \alpha_k C^{-\frac{\varepsilon_k}{\theta(C)-1}}. \tag{A.6}$$

Hence,

$$\mathcal{H}(C; \boldsymbol{c}) = \left[\sum_{k=1}^{K} \alpha_k \int_{\mathcal{I}_k} \left(\frac{c_{kj}}{K_k(C)} \right)^{\rho(C)} dj \right]^{\frac{1}{\rho(C)}}. \tag{A.7}$$

Because

$$\sum_{k=1}^{K} \alpha_k = 1$$

and every interval $\mathcal{I}_k$ has unit measure, (A.7) is a properly normalised weighted power mean of the quantities $c_{kj}/K_k(C)$.

Step 3: $\mathcal{H}(C;\boldsymbol{c})$ is strictly decreasing in C. From (A.2), $\theta'(C) = \delta_1\theta(C)/C < 0$. Consequently, $\rho'(C) = \theta'(C)/\theta(C)^2 < 0$.

Next, differentiating (A.6) gives

$$\frac{K_k'(C)}{K_k(C)} = -\frac{\varepsilon_k}{C[\theta(C)-1]} + \frac{\varepsilon_k\theta'(C)\log C}{[\theta(C)-1]^2}. \tag{A.8}$$

The first term in (A.8) is strictly positive because $\varepsilon_k < 0,\ \theta(C) > 1$. The second term is nonnegative because $\varepsilon_k < 0,\ \theta'(C) < 0,$ and $\log C \geq 0$ on $\mathcal{D} = [1, \overline{C}]$. Therefore,

$$K_k'(C) > 0 \qquad \text{for every } C \in \mathcal{D} \text{ and every } k. \tag{A.9}$$

Consider any $C_2 > C_1$ in $\mathcal{D}$. Equation (A.9) implies

$$\frac{c_{kj}}{K_k(C_2)} < \frac{c_{kj}}{K_k(C_1)} \qquad \text{for every } k, j.$$

We also know that $\rho(C_2) < \rho(C_1)$. Let $M_r(\boldsymbol{x})$ denote the normalised weighted power mean of order r. For positive arguments, the power mean is strictly increasing in each argument and weakly increasing in its order r. It follows that

$$\begin{aligned}
\mathcal{H}(C_2;\boldsymbol{c}) &= M_{\rho(C_2)}\left(\left\{\frac{c_{kj}}{K_k(C_2)}\right\}_{k,j}\right) \\
&< M_{\rho(C_2)}\left(\left\{\frac{c_{kj}}{K_k(C_1)}\right\}_{k,j}\right) \\
&\leq M_{\rho(C_1)}\left(\left\{\frac{c_{kj}}{K_k(C_1)}\right\}_{k,j}\right) \\
&= \mathcal{H}(C_1;\boldsymbol{c}).
\end{aligned}$$

The first inequality is strict because all normalised quantities strictly decline when C rises. Thus, $\mathcal{H}(C;\boldsymbol{c})$ is continuous and strictly decreasing in C.

Step 4: Existence and uniqueness. Fix any $\boldsymbol{c} \in \mathfrak{B}_0$. By the boundary-crossing condition,

$$\mathcal{H}(1;\boldsymbol{c}) > 1 > \mathcal{H}(\overline{C};\boldsymbol{c}).$$

Because $\mathcal{H}(C;\boldsymbol{c})$ is continuous in C, the intermediate value theorem implies that there exists some $C(\boldsymbol{c}) \in (1, \overline{C})$ such that $\mathcal{H}(C(\boldsymbol{c});\boldsymbol{c}) = 1$. Because $\mathcal{H}(C;\boldsymbol{c})$ is strictly decreasing in C, this solution is unique. By Step 1, it is exactly the unique solution to (A.5).

Step 5: Strict monotonicity and quasi-concavity. First, consider two bundles $\boldsymbol{c}, \boldsymbol{c}' \in \mathfrak{B}_0$ satisfying

$$c'_{kj} \geq c_{kj} \qquad \text{for almost every } (k, j),$$

with strict inequality on a set of positive measure.

For every fixed $C \in \mathcal{D}$, equation (A.7) and the strict monotonicity of the power mean in its arguments imply

$$\mathcal{H}(C;\boldsymbol{c}') > \mathcal{H}(C;\boldsymbol{c}).$$

In particular,

$$\mathcal{H}(C(\boldsymbol{c});\boldsymbol{c}') > \mathcal{H}(C(\boldsymbol{c});\boldsymbol{c}) = 1.$$

Because $\mathcal{H}(C;\boldsymbol{c}')$ is strictly decreasing in C, its unique crossing of one must occur at a higher value: $C(\boldsymbol{c}') > C(\boldsymbol{c})$. Thus, $C(\boldsymbol{c})$ is strictly increasing.

To establish quasi-concavity, fix an arbitrary convex set $\mathfrak{B} \subseteq \mathfrak{B}_0$ and an arbitrary $c \in \mathcal{D}$. Since $\mathcal{H}(C;\boldsymbol{c})$ is strictly decreasing in its first argument,

$$C(\boldsymbol{c}) \geq c \quad \Longleftrightarrow \quad \mathcal{H}(c;\boldsymbol{c}) \geq 1.$$

For fixed c, the quantities $\rho(c) \in (0,1)$ and $K_k(c) > 0$ are constants, and

$$\mathcal{H}(c;\boldsymbol{c}) = \left[\sum_{k=1}^{K} \alpha_k \int_{\mathcal{I}_k} \left(\frac{c_{kj}}{K_k(c)}\right)^{\rho(c)} dj\right]^{\frac{1}{\rho(c)}}. \tag{A.10}$$

Because $\rho(c) \in (0,1)$, equation (A.10) is a concave CES aggregator in the consumption bundle $\boldsymbol{c}$. Consequently, its upper contour set

$$\{\boldsymbol{c} : \mathcal{H}(c;\boldsymbol{c}) \geq 1\}$$

is convex. Since $\mathfrak{B}$ is convex, the set

$$\{\boldsymbol{c} \in \mathfrak{B} : C(\boldsymbol{c}) \geq c\} = \mathfrak{B} \cap \{\boldsymbol{c} : \mathcal{H}(c;\boldsymbol{c}) \geq 1\}$$

is convex. Hence, every upper contour set of $C(\boldsymbol{c})$ on $\mathfrak{B}$ is convex, so the implicit consumption aggregator is quasi-concave.

This completes the proof. □

A.2 Deriving the Hicksian Demand

It is straightforward to show that given prices and a sectoral composite C_{kt}, households optimally choose their allocation among $\{C_{k\iota t}\}$, resulting in the following demand function:

$$C_{k\iota t} = \left(\frac{P_{k\iota t}}{P_{kt}}\right)^{-\theta_t} C_{kt},$$

which is equation (21) in the main text.

Given this result, we then consider household h's utility maximisation problem on the allocation of $\{C_{kt}\}$, formulated as follows:

$$\max_{\{C_{kt}\}} C_t \tag{A.11}$$

subject to

$$1 = \sum_{k=1}^{K} [g_k(C_t)]^{\frac{1}{\theta_t}} (C_{kt})^{1-\frac{1}{\theta_t}}, \tag{A.12}$$

$$\sum_{k=1}^{K} P_{kt} C_{kt} \leq E_t. \tag{A.13}$$

We use ψ and λ to denote the Lagrange multipliers on constraint (A.12) and budget constraint (A.13), respectively:

$$\mathcal{L} = C_t + \psi \left(\sum_{k=1}^{K} [g_k(C_t)]^{\frac{1}{\theta_t}} (C_{kt})^{1-\frac{1}{\theta_t}} - 1 \right) + \lambda \left(E_t - \sum_{k=1}^{K} P_{kt} C_{kt} \right)$$

The first order condition w.r.t. C_{kt} is given by

$$\psi \frac{\theta_t - 1}{\theta_t} m_{kt} = \lambda P_{kt} C_{kt}, \tag{A.14}$$

where we define $m_{kt} = C_{kt}^{1-1/\theta_t} [g_k(C_t)]^{1/\theta_t}$. It follows directly from constraint (A.12) that $\sum m_{kt} = 1$. Note from equation (A.14) that expenditure $P_{kt} C_{kt}$ is proportional to m_{kt}, so that m_{kt} is the expenditure share on quality group k. Summing (A.14) across quality groups we have $E = (\psi/\lambda)(\theta_t - 1)/\theta_t$. Substituting the expression of m_{kt} and E_t into equation (A.14) yields the demand function $C_{kt} = g_k(C_t)(P_{kt}/E_t)^{-\theta_t}$. Substituting these demand functions into the budget constraint yields the expenditure function

$$\mathcal{E}(C_t; \boldsymbol{P}_t) = \left(\sum_{k=1}^{K} g_k(C_t) (P_{kt})^{1-\theta_t} \right)^{\frac{1}{1-\theta_t}}.$$

A.3 Implicit Marshallian Demand

In Section 4.2, we derive households' Hicksian demand. Here, we characterise the corresponding Marshallian demand.

By Lemma 1, for any strictly positive price vector $\boldsymbol{P}_t$, the expenditure function $\mathcal{E}^h(C; \boldsymbol{P}_t)$ is strictly increasing in C over the admissible domain $\mathcal{D} = [1, \overline{C}]$. Hence, for every expenditure level $E_t^h \in \mathcal{E}^h(\mathcal{D}; \boldsymbol{P}_t)$, there exists a unique real-consumption level

$$C_t^h = C^h(E_t^h; \boldsymbol{P}_t), \tag{A.15}$$

where $C^h(\cdot; \boldsymbol{P}_t)$ denotes the inverse of the expenditure function.

The corresponding elasticity of substitution is therefore uniquely determined by

$$\theta_t^h = \exp\left[\delta_0 + \delta_1 \log C^h(E_t^h; \boldsymbol{P}_t)\right]. \tag{A.16}$$

Substituting equations (A.15) and (A.16) into the Hicksian demand functions gives the Marshallian demand for variety ι in quality group k:

$$C_{k\iota t}^h = \alpha_k \left[C^h(E_t^h; \boldsymbol{P}_t)\right]^{\varepsilon_k} \left(\frac{P_{k\iota t}}{E_t^h}\right)^{-\theta_t^h}.$$

The corresponding expenditure share of variety ι is

$$m_{k\iota t}^h \equiv \frac{P_{k\iota t} C_{k\iota t}^h}{E_t^h} = \alpha_k \left[C^h(E_t^h; \boldsymbol{P}_t)\right]^{\varepsilon_k} \left(\frac{P_{k\iota t}}{E_t^h}\right)^{1-\theta_t^h}.$$

Integrating over varieties in quality group k gives the group-level expenditure share:

$$m_{kt}^h = \alpha_k \left[C^h(E_t^h; \boldsymbol{P}_t)\right]^{\varepsilon_k} \int_{k-1}^{k} \left(\frac{P_{k\iota t}}{E_t^h}\right)^{1-\theta_t^h} d\iota.$$

Thus, Marshallian demand and expenditure shares are uniquely determined by total expenditure and the price vector:

$$C_{k\iota t}^h = \mathcal{D}_{k\iota}^h(E_t^h, \boldsymbol{P}_t), \qquad m_{kt}^h = \mathcal{M}_k^h(E_t^h, \boldsymbol{P}_t).$$

A.4 Savers' Optimisation Problem

The following proposition characterises the savers' optimisation problem. In particular, we demonstrate that savers' optimisation problem can be decomposed into two separate sub-problems, and we provide sufficient conditions for the solution to be unique.

For simplicity, we temporarily drop the saver superscript S. Define the elasticity of the expenditure function

with respect to real consumption and the elasticity of that elasticity with respect to real consumption as

$$\eta^{C}_{\mathcal{E},t} \equiv \frac{\partial \log \mathcal{E}_t(C_t; \boldsymbol{P}_t)}{\partial \log C_t}, \qquad \eta^{C}_{\eta^{C}_{\mathcal{E}},t} \equiv \frac{\partial \log \eta^{C}_{\mathcal{E},t}}{\partial \log C_t}.$$

Proposition A.1. *Consider a saver's optimisation problem with preferences specified by equations (9), (10), (17) and (18). Suppose conditions in Lemma A.1 and the following condition hold,*

$$\eta^{C}_{\mathcal{E},t} + \eta^{C}_{\eta^{C}_{\mathcal{E}},t} > 1 - \frac{1}{\vartheta}. \tag{A.17}$$

Then given the sequences $\{W_t, \boldsymbol{P}_t, \phi_t, i_t, D_t, T_t\}_{t=0}^{\infty}$ and initial bond holdings B_0, the saver's problem has a unique solution. The solution can be characterised as follows.

1. *intratemporal conditions: i) Given total expenditure E_t, the intratemporal allocations of consumption goods $\{\boldsymbol{C}_t; E_t = \sum P_{kt} C_{kt}\}$ satisfy conditions (21), (22) and the expenditure function relating E_t and C_t, and ii) the labour supply satisfies the following condition:*

$$\phi_t \frac{(C_t)^{-1/\vartheta}}{(N_t)^{1/\gamma}} = \frac{\partial \mathcal{E}(C_t; \boldsymbol{P}_t)/\partial C_t}{oW_t} \tag{A.18}$$

2. *intertemporal conditions: the intertemporal allocation of expenditures $\{E_t\}$ satisfies the Euler equation:*

$$\phi_t C_t^{-1/\vartheta} \frac{\partial \mathcal{C}_t(E_t; \boldsymbol{P}_t)}{\partial E_t} = \beta(1 + i_t) \mathrm{E}_t \left[\phi_{t+1} C_{t+1}^{-1/\vartheta} \frac{\partial \mathcal{C}_{t+1}(E_{t+1}; \boldsymbol{P}_{t+1})}{\partial E_{t+1}} \right], \tag{A.19}$$

   together with the period budget constraints and the transversality condition

$$\lim_{T \to \infty} \mathrm{E}_0 \left[\beta^T \mu_T \frac{B_{T+1}}{1 + i_T} \right] = 0, \quad \textit{where} \quad \mu_t \equiv \phi_t C_t^{-1/\vartheta} \frac{\partial \mathcal{C}_t(E_t; \boldsymbol{P}_t)}{\partial E_t} = \mathcal{V}_t'(x_t). \tag{A.20}$$

Proof. The proof proceeds in five steps.

Step 1: Reduction of the differentiated-goods problem. For any given nominal expenditure E_t, the household chooses the differentiated consumption bundle to maximise real aggregate consumption:

$$\mathcal{C}_t(E_t; \boldsymbol{P}_t) = \max_{\boldsymbol{C}_t \geq 0} \left\{ \mathcal{C}(\boldsymbol{C}_t) : \sum_{k=1}^{K} \int_{k-1}^{k} P_{k\iota t} C_{k\iota t} \, d\iota \leq E_t \right\}.$$

Equivalently, the associated expenditure-minimisation problem is

$$\mathcal{E}_t(C_t; \boldsymbol{P}_t) = \min_{\boldsymbol{C}_t \geq 0} \left\{ \sum_{k=1}^{K} \int_{k-1}^{k} P_{k\iota t} C_{k\iota t} \, d\iota : \mathcal{C}(\boldsymbol{C}_t) \geq C_t \right\}.$$

By Lemma A.1 and Lemma 1, $\mathcal{E}_t(C_t; \boldsymbol{P}_t)$ is strictly increasing in C_t. Its inverse $C_t = C_t(E_t; \boldsymbol{P}_t)$ is therefore well defined and strictly increasing in E_t. Consequently, after solving the within-period allocation across goods, the consumption part of period utility can be written as a function of total expenditure:

$$u_t(E_t) \equiv \frac{\phi_t C_t(E_t; \boldsymbol{P}_t)^{1-1/\vartheta}}{1-1/\vartheta}.$$

Step 2: Concavity of consumption subutility in expenditure. To simplify notation, write

$$C_t = C_t(E_t; \boldsymbol{P}_t), \qquad \eta_t = \eta^C_{\mathcal{E},t}.$$

The marginal utility of expenditure is given by

$$u_t'(E_t) = \phi_t C_t^{-1/\vartheta} \frac{\partial C_t}{\partial E_t} = \frac{\phi_t C_t^{1-1/\vartheta}}{\eta_t E_t} > 0. \tag{A.21}$$

Taking the logarithmic derivative of (A.21) with respect to E_t gives

$$\begin{aligned} \frac{\partial \log u_t'(E_t)}{\partial \log E_t} &= \left(1 - \frac{1}{\vartheta}\right) \frac{\partial \log C_t}{\partial \log E_t} - \frac{\partial \log \eta_t}{\partial \log E_t} - 1 \\ &= \frac{1 - 1/\vartheta - \eta^C_{\mathcal{E},t} - \eta^C_{\eta^C_{\mathcal{E}},t}}{\eta^C_{\mathcal{E},t}}. \end{aligned}$$

Condition (A.17) implies

$$\frac{\partial \log u_t'(E_t)}{\partial \log E_t} < 0.$$

Since $u_t'(E_t) > 0$, it follows that

$$u_t''(E_t) < 0.$$

Thus, consumption subutility is strictly increasing and strictly concave in nominal expenditure.

Step 3: Static expenditure and labour problem. Let $w_t \equiv oW_t > 0$ denote the household's effective nominal wage. For a given amount x_t of net nonlabour-resource use, consider

$$\mathcal{V}_t(x_t) = \max_{E_t, N_t \geq 0} \left\{ u_t(E_t) - \frac{N_t^{1+1/\gamma}}{1+1/\gamma} : E_t - w_t N_t \leq x_t \right\}. \tag{A.22}$$

The objective function in (A.22) is strictly concave in (E_t, N_t): $u_t(E_t)$ is strictly concave by Step 2, while $-N_t^{1+1/\gamma}/(1+1/\gamma)$ is strictly concave for $\gamma > 0$. The feasible set is convex because its constraint is affine. Hence, the static problem has at most one solution. Existence follows from the maintained regularity and feasibility assumptions.

Combining the budget constraint $x_t + w_t N_t - E_t \geq 0$ with the first-order conditions gives labour supply as

a function of expenditure:

$$N_t(E_t) = \left[w_t u_t'(E_t)\right]^{\gamma} \tag{A.23}$$

Thus, labour is determined jointly with expenditure by (A.23) and

$$E_t - w_t N_t(E_t) = x_t.$$

Equivalently, we obtain the intratemporal labour-supply condition as (A.18).

Step 4: Concavity of indirect period utility. We next establish that $\mathcal{V}_t(x_t)$ is strictly concave. Consider $x_t^1 \neq x_t^2$, and let (E_t^1, N_t^1) and (E_t^2, N_t^2) be their respective unique optimisers. For any $\lambda \in (0,1)$, define $x_t^\lambda = \lambda x_t^1 + (1-\lambda)x_t^2$ and $(E_t^\lambda, N_t^\lambda) = \lambda(E_t^1, N_t^1) + (1-\lambda)(E_t^2, N_t^2)$. Because the constraint is affine, $(E_t^\lambda, N_t^\lambda)$ is feasible for x_t^λ. Hence,

$$\mathcal{V}_t(x_t^\lambda) \geq u_t(E_t^\lambda) - \frac{(N_t^\lambda)^{1+1/\gamma}}{1+1/\gamma}.$$

The objective is strictly concave in (E_t, N_t). Moreover, $(E_t^1, N_t^1) \neq (E_t^2, N_t^2)$, because both resource constraints bind and $x_t^1 \neq x_t^2$. Therefore,

$$\begin{aligned}\mathcal{V}_t(x_t^\lambda) &> \lambda\left[u_t(E_t^1) - \frac{(N_t^1)^{1+1/\gamma}}{1+1/\gamma}\right] + (1-\lambda)\left[u_t(E_t^2) - \frac{(N_t^2)^{1+1/\gamma}}{1+1/\gamma}\right] \\ &= \lambda\mathcal{V}_t(x_t^1) + (1-\lambda)\mathcal{V}_t(x_t^2).\end{aligned}$$

Hence, $\mathcal{V}_t$ is strictly concave.

By the envelope theorem,

$$\mathcal{V}_t'(x_t) = \mu_t = u_t'(E_t) = \phi_t C_t^{-1/\vartheta} \frac{\partial \boldsymbol{C}_t(E_t; \boldsymbol{P}_t)}{\partial E_t} > 0. \tag{A.24}$$

Thus, $\mathcal{V}_t$ is also strictly increasing.

Step 5: Reduction to the bond-only problem. Using $x_t = E_t - w_t N_t$, the original period budget constraint becomes

$$x_t + \frac{B_{t+1}}{1+i_t} = B_t + \frac{D_t}{\lambda_S} + T_t.$$

Therefore, after solving the static expenditure-and-labour problem, the saver chooses only bond holdings:

$$\max_{\{B_{t+1}\}_{t=0}^{\infty}} \quad \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t \mathcal{V}_t\left(B_t + \frac{D_t}{\lambda_S} + T_t - \frac{B_{t+1}}{1+i_t}\right).$$

Because $\mathcal{V}_t$ is strictly concave and the mapping from any feasible bond path $\{B_{t+1}\}_{t\geq 0}$, given B_0, to the implied resource-use path $\{x_t\}_{t\geq 0}$ is affine, the lifetime objective is strictly concave over the convex feasible set. Hence, we can then apply results from discrete dynamic programming (e.g. Acemoglu (2008) Theorem 6.12 or Miao (2020) Theorem 8.6.1 and Theorem 8.6.2) to ensure that the following conditions provide sufficient and necessary conditions for a sequence of $\{B_{t+1}\}_t$ to characterise the unique solution of savers' optimisation problem.

The first-order condition for B_{t+1} gives

$$\mathcal{V}_t'(x_t) = \beta(1+i_t)\mathrm{E}_t\left[\mathcal{V}_{t+1}'(x_{t+1})\right]. \tag{A.25}$$

Substituting (A.24) into (A.25) yields

$$\phi_t C_t^{-1/\vartheta}\frac{\partial \mathcal{C}_t(E_t;\boldsymbol{P}_t)}{\partial E_t} = \beta(1+i_t)\mathrm{E}_t\left[\phi_{t+1}C_{t+1}^{-1/\vartheta}\frac{\partial \mathcal{C}_{t+1}(E_{t+1};\boldsymbol{P}_{t+1})}{\partial E_{t+1}}\right].$$

This is the saver Euler equation.

Finally, imposing the transversality condition

$$\lim_{T\to\infty}\mathrm{E}_0\left[\beta^T\mathcal{V}_T'(x_T)\frac{B_{T+1}}{1+i_T}\right] = 0$$

rules out explosive asset-accumulation paths. The Euler equation, the static expenditure-and-labour conditions, the period budget constraints, and the transversality condition therefore characterise the unique solution. This completes the proof. □

B. Proof of Theorems and Propositions

Notation Summary. Table B.1 summarises variables and notations in our paper.

Table B.1: **Notation Summary**

Households	
$h \in \{H, S\}$	Household type: H = hand-to-mouth, S = savers
C_t^h	Real consumption of type-h household
E_t^h	Nominal expenditure of type-h household
P_t^h	Household-specific price index
N_t^h	Labour supply of type-h household
ϕ_t^h	Preference (demand) shifter for type-h household
o^h	Labour efficiency of type-h household
$C^h(E_t^h; \mathbf{P}_t)$	Indirect utility / real consumption function
Goods, Quality, and Prices	
k	Quality-group index
P_{it}	Nominal price of product i
p_{it}	Log nominal price of product i
P_{kt}	Aggregate price index for quality group k
p_{kt}	Log price level for quality group k
P_t	Aggregate price level
π_{it}	Inflation rate of product i
π_{kt}	Inflation rate of quality group k
π_t	Aggregate inflation
Firms and Markups	
μ_{it}	Actual markup of firm (log deviation) i
M_{kt}^*	Ideal (flexible-price) markup for quality group k
μ_{kt}^*	Ideal (flexible-price) markup for quality group k (log deviation)
MC_t	Nominal marginal cost
w_t	Log nominal wage
w_{kt}	Real wage relative to quality-group price, $w_t - p_{kt}$
ζ_k	Calvo probability of price adjustment in group k
κ_k	Slope coefficient of quality-group Phillips curve

Income, Output, and Aggregates	
Y_t	GDP
Y_{kt}	Output of quality group k
W_t	Nominal wage level
i_t	Net nominal interest rate
r_t	Net real interest rate

Preferences and Policy Parameters	
β	Household discount factor
ϑ	Intertemporal elasticity parameter in utility
γ	Frisch elasticity parameter
π^*	Inflation target
φ_π	Taylor rule coefficient on inflation
ε_k	Income-elasticity parameter

Other Objects	
Δx_t	First difference of variable x_t
E_t	Expectations conditional on information at time t

B.1 Proof of Lemma 1

For simplicity, suppress the household superscript h and the time subscript t. Let $\boldsymbol{c} = \{c_{k\iota}\}_{k,\iota}$ denote a strictly positive consumption bundle. Recall the function $\mathcal{H}(C;\boldsymbol{c})$ defined in (A.7). Define the natural admissible domain

$$\mathfrak{B}_0 \equiv \left\{\boldsymbol{c} \gg 0 : \mathcal{H}(1;\boldsymbol{c}) > 1 > \mathcal{H}(\overline{C};\boldsymbol{c})\right\}.$$

By Lemma A.1, for every $\boldsymbol{c} \in \mathfrak{B}_0$, there exists a unique $C(\boldsymbol{c}) \in (1, \overline{C})$ satisfying $\mathcal{H}(C(\boldsymbol{c});\boldsymbol{c}) = 1$. Moreover, because $\mathcal{H}(C;\boldsymbol{c})$ is strictly decreasing in C, for every $C \in (1, \overline{C})$,

$$C(\boldsymbol{c}) \geq C \quad \Longleftrightarrow \quad \mathcal{H}(C;\boldsymbol{c}) \geq 1.$$

Consequently, the expenditure function can equivalently be written as

$$\mathcal{E}(C;\boldsymbol{P}) = \min_{\boldsymbol{c}\in\mathfrak{B}_0} \left\{ \sum_{k=1}^{K} \int_{\mathcal{I}_k} P_{k\iota} c_{k\iota}\, d\iota : \mathcal{H}(C;\boldsymbol{c}) \geq 1 \right\}. \tag{B.1}$$

The strictly positive Hicksian bundle derived in Appendix A.2 attains this minimum when prices are strictly positive and the corresponding sectoral price indices are finite.

For any fixed C, equation (A.7) implies that $\mathcal{H}(C;\boldsymbol{c})$ is homogeneous of degree one in $\boldsymbol{c}$:

$$\mathcal{H}(C;\lambda\boldsymbol{c}) = \lambda\mathcal{H}(C;\boldsymbol{c}) \qquad \text{for every } \lambda > 0.$$

Consider two target consumption levels $1 < C_1 < C_2 < \overline{C}$. Let $\boldsymbol{c}_2^*$ be a cost-minimising bundle for target C_2. The constraint in (B.1) must bind:

$$\mathcal{H}(C_2;\boldsymbol{c}_2^*) = 1. \tag{B.2}$$

Indeed, suppose instead that $\mathcal{H}(C_2;\boldsymbol{c}_2^*) > 1$. Define

$$\eta \equiv \frac{1}{\mathcal{H}(C_2;\boldsymbol{c}_2^*)} \in (0,1).$$

By homogeneity, $\mathcal{H}(C_2;\eta\boldsymbol{c}_2^*) = 1$, while strict positivity of prices implies

$$\sum_{k=1}^{K}\int_{\mathcal{I}_k} P_{k\iota}\eta c_{2,k\iota}^*\,d\iota < \sum_{k=1}^{K}\int_{\mathcal{I}_k} P_{k\iota}c_{2,k\iota}^*\,d\iota,$$

contradicting the optimality of $\boldsymbol{c}_2^*$.

Because $\mathcal{H}(C;\boldsymbol{c})$ is strictly decreasing in its first argument and $C_1 < C_2$, equation (B.2) implies

$$q \equiv \mathcal{H}(C_1;\boldsymbol{c}_2^*) > \mathcal{H}(C_2;\boldsymbol{c}_2^*) = 1.$$

Define $\lambda \equiv 1/q \in (0,1)$. Using homogeneity once more gives

$$\mathcal{H}(C_1;\lambda\boldsymbol{c}_2^*) = \lambda\mathcal{H}(C_1;\boldsymbol{c}_2^*) = 1.$$

Since $1 < C_1 < \overline{C}$ and $\mathcal{H}(C;\boldsymbol{c})$ is strictly decreasing in C, this equality also implies

$$\mathcal{H}(1;\lambda\boldsymbol{c}_2^*) > 1 > \mathcal{H}(\overline{C};\lambda\boldsymbol{c}_2^*).$$

Thus, $\lambda\boldsymbol{c}_2^* \in \mathfrak{B}_0$, and it is feasible for the expenditure-minimisation problem associated with target C_1. Therefore,

$$\begin{aligned}\mathcal{E}(C_1;\boldsymbol{P}) &\le \sum_{k=1}^{K}\int_{\mathcal{I}_k} P_{k\iota}\lambda c_{2,k\iota}^*\,d\iota = \lambda\sum_{k=1}^{K}\int_{\mathcal{I}_k} P_{k\iota}c_{2,k\iota}^*\,d\iota \\ &= \lambda\mathcal{E}(C_2;\boldsymbol{P}) < \mathcal{E}(C_2;\boldsymbol{P}).\end{aligned}$$

Hence,

$$\mathcal{E}(C_1;\boldsymbol{P}) < \mathcal{E}(C_2;\boldsymbol{P}) \qquad \text{whenever } C_1 < C_2,$$

so the expenditure function is strictly increasing in the target consumption level.

B.2 Proof of Proposition 1

For notational simplicity, suppress the time subscript and define $\theta_H \equiv \theta_t^H, \qquad \theta_S \equiv \theta_t^S$. Since $\delta_1 < 0$ and $C_t^H < C_t^S$, we have $\theta_H > \theta_S > 1$. Consequently, $\underline{M} \equiv \theta_H/(\theta_H - 1) < \theta_S/(\theta_S - 1) \equiv \overline{M}$.

Define the relative demand of hand-to-mouth households and savers as

$$R_k(M) \equiv \frac{C_{kt}^H}{C_{kt}^S} = A_k(W_t M)^{\theta_S - \theta_H},$$

where

$$A_k \equiv \left(\frac{C_t^H}{C_t^S}\right)^{\varepsilon_k} \frac{(E_t^H)^{\theta_H}}{(E_t^S)^{\theta_S}} > 0.$$

Because $\theta_S - \theta_H < 0$, relative demand is decreasing in the markup:

$$R_k'(M) = (\theta_S - \theta_H)\frac{R_k(M)}{M} < 0. \tag{B.3}$$

Rearranging the first-order condition of profit-maximisation for $P_{\iota t}$ gives

$$P_{kt} \sum_{h \in \{H,S\}} \lambda_h(\theta_t^h - 1)C_{kt}^h = W_t \sum_{h \in \{H,S\}} \lambda_h \theta_t^h C_{kt}^h. \tag{B.4}$$

Given $R_k(M)$, define the markup implied by the firm's first-order condition as

$$\mathcal{M}_k(M) \equiv \frac{\lambda_H \theta_H R_k(M) + \lambda_S \theta_S}{\lambda_H(\theta_H - 1)R_k(M) + \lambda_S(\theta_S - 1)}. \tag{B.5}$$

The equilibrium ideal markup is a fixed point of this mapping:

$$M_{kt}^* = \mathcal{M}_k(M_{kt}^*).$$

Equivalently, define the root function

$$\Phi_k(M) \equiv \mathcal{M}_k(M) - M.$$

The ideal markup M_{kt}^* is characterised by

$$\Phi_k(M_{kt}^*) = 0.$$

Existence. For any $R > 0$, define

$$\mathcal{M}(R) = \frac{\lambda_H \theta_H R + \lambda_S \theta_S}{\lambda_H(\theta_H - 1)R + \lambda_S(\theta_S - 1)}.$$

First,

$$\mathcal{M}(R) - \frac{\theta_H}{\theta_H - 1} = \frac{\lambda_S(\theta_H - \theta_S)}{(\theta_H - 1)\left[\lambda_H(\theta_H - 1)R + \lambda_S(\theta_S - 1)\right]} > 0.$$

Similarly,

$$\mathcal{M}(R) - \frac{\theta_S}{\theta_S - 1} = -\frac{\lambda_H R(\theta_H - \theta_S)}{(\theta_S - 1)\left[\lambda_H(\theta_H - 1)R + \lambda_S(\theta_S - 1)\right]} < 0.$$

Therefore,

$$\underline{M} < \mathcal{M}_k(M) < \overline{M} \qquad \text{for every } M > 0.$$

It follows that

$$\Phi_k(\underline{M}) > 0, \qquad \Phi_k(\overline{M}) < 0.$$

Since $\Phi_k(M)$ is continuous, the intermediate value theorem implies that there exists at least one

$$M_k^* \in (\underline{M}, \overline{M})$$

such that $\Phi_k(M_k^*) = 0$.

Derivative and uniqueness. Let

$$D_k(M) \equiv \lambda_H(\theta_H - 1)R_k(M) + \lambda_S(\theta_S - 1).$$

Differentiating the markup mapping in (B.5) with respect to R_k gives

$$\begin{aligned} \frac{\partial \mathcal{M}_k}{\partial R_k} &= \frac{\lambda_H \theta_H D_k - \lambda_H(\theta_H - 1)\left[\lambda_H \theta_H R_k + \lambda_S \theta_S\right]}{D_k^2} \\ &= \frac{\lambda_H \lambda_S(\theta_S - \theta_H)}{D_k^2}. \end{aligned} \tag{B.6}$$

Combining (B.3) and (B.6) yields

$$\begin{aligned} \mathcal{M}_k'(M) &= \frac{\partial \mathcal{M}_k}{\partial R_k} R_k'(M) \\ &= \frac{\lambda_H \lambda_S(\theta_H - \theta_S)^2 R_k(M)}{M D_k(M)^2}. \end{aligned}$$

Hence,

$$\Phi_k'(M) = \frac{\lambda_H \lambda_S(\theta_H - \theta_S)^2 R_k(M)}{M\left[\lambda_H(\theta_H - 1)R_k(M) + \lambda_S(\theta_S - 1)\right]^2} - 1.$$

A sufficient condition for uniqueness is

$$\sup_{M\in[\underline{M},\overline{M}]} \frac{\lambda_H\lambda_S(\theta_H-\theta_S)^2 R_k(M)}{M\left[\lambda_H(\theta_H-1)R_k(M)+\lambda_S(\theta_S-1)\right]^2} < 1. \tag{B.7}$$

Under (B.7), $\Phi_k'(M) < 0$ throughout $[\underline{M}, \overline{M}]$. Since Φ_k is strictly decreasing and changes sign exactly once, the root M_k^* is unique.

A primitive sufficient condition. A simpler, although more conservative, sufficient condition can be obtained by noting that

$$D_k(M) \geq \lambda_S(\theta_S - 1).$$

It follows that

$$\begin{aligned}\mathcal{M}_k'(M) &\leq \frac{\lambda_H}{\lambda_S}\frac{(\theta_H-\theta_S)^2}{(\theta_S-1)^2}\frac{R_k(M)}{M} \\ &= \frac{\lambda_H}{\lambda_S}\frac{(\theta_H-\theta_S)^2}{(\theta_S-1)^2}A_k W_t^{\theta_S-\theta_H} M^{\theta_S-\theta_H-1}.\end{aligned} \tag{B.8}$$

Since

$$\theta_S - \theta_H - 1 < 0$$

and $M \geq \underline{M}$, the right-hand side of (B.8) is bounded above by its value at $\underline{M}$. Therefore, the following is sufficient for $\Phi_k'(M) < 0$ throughout the relevant interval:

$$\begin{aligned}&\frac{\lambda_H}{\lambda_S}\frac{(\theta_H-\theta_S)^2}{(\theta_S-1)^2} W_t^{\theta_S-\theta_H}\left(\frac{\theta_H}{\theta_H-1}\right)^{\theta_S-\theta_H-1} \\ &\quad\times\left(\frac{C_t^H}{C_t^S}\right)^{\varepsilon_k}\frac{(E_t^H)^{\theta_H}}{(E_t^S)^{\theta_S}} < 1.\end{aligned} \tag{B.9}$$

This establishes uniqueness under the stated sufficient condition.

Finally, because $\theta_H > 1$ and $\theta_S > 1$, nominal profits are zero at $M = 1$, strictly positive for $M > 1$ sufficiently close to one, and converge to zero as $M \to \infty$. Hence, a global maximiser exists in the interior of $M > 1$ and must satisfy the firm's first-order condition. Since the first-order condition has a unique solution under (38), this solution is the firm's unique global profit-maximising ideal markup.

B.3 Proof of Proposition 2

Continue to define

$$R_k(M, \varepsilon_k) = \left(\frac{C_t^H}{C_t^S}\right)^{\varepsilon_k}\frac{(E_t^H)^{\theta_H}}{(E_t^S)^{\theta_S}}(W_t M)^{\theta_S-\theta_H},$$

and let

$$\Phi_k(M, \varepsilon_k) = \mathcal{M}_k(M, \varepsilon_k) - M.$$

The ideal markup is the unique root satisfying

$$\Phi_k(M_k^*, \varepsilon_k) = 0,$$

or, equivalently,

$$M_k^* = \mathcal{M}_k(M_k^*, \varepsilon_k).$$

Because $C_t^H / C_t^S < 1$,

$$\frac{\partial R_k}{\partial \varepsilon_k} = R_k \log\left(\frac{C_t^H}{C_t^S}\right) < 0.$$

Moreover, from (B.6),

$$\frac{\partial \mathcal{M}_k}{\partial R_k} = \frac{\lambda_H \lambda_S (\theta_S - \theta_H)}{[\lambda_H(\theta_H - 1)R_k + \lambda_S(\theta_S - 1)]^2} < 0.$$

Therefore,

$$\frac{\partial \Phi_k}{\partial \varepsilon_k} = \frac{\partial \mathcal{M}_k}{\partial R_k} \frac{\partial R_k}{\partial \varepsilon_k} > 0.$$

Under the uniqueness condition in Proposition 1, $\partial \Phi_k / \partial M < 0$. The implicit function theorem therefore gives

$$\frac{dM_k^*}{d\varepsilon_k} = -\frac{\partial \Phi_k / \partial \varepsilon_k}{\partial \Phi_k / \partial M} > 0.$$

Hence, if $\varepsilon_k < \varepsilon_j$, then $M_k^* < M_j^*$.

It remains to compare customer composition. Evaluating relative demand at the equilibrium markup gives

$$R_k^* = R_k(M_k^*, \varepsilon_k).$$

Its total derivative satisfies

$$\frac{d \log R_k^*}{d\varepsilon_k} = \log\left(\frac{C_t^H}{C_t^S}\right) + (\theta_S - \theta_H)\frac{1}{M_k^*}\frac{dM_k^*}{d\varepsilon_k} < 0,$$

because

$$\log\left(\frac{C_t^H}{C_t^S}\right) < 0, \qquad \theta_S - \theta_H < 0, \qquad \frac{dM_k^*}{d\varepsilon_k} > 0.$$

Thus, a lower ε_k implies a larger equilibrium relative demand R_k^*.

Finally, the revenue share attributable to hand-to-mouth households is

$$\omega_{kt}^{H} = \frac{\lambda_H R_k^*}{\lambda_H R_k^* + \lambda_S}.$$

Since

$$\frac{\partial \omega_{kt}^{H}}{\partial R_k^*} = \frac{\lambda_H \lambda_S}{(\lambda_H R_k^* + \lambda_S)^2} > 0,$$

a lower ε_k implies a larger ω_{kt}^{H}. Therefore, if $\varepsilon_k < \varepsilon_j$, then

$$\omega_{kt}^{H} > \omega_{jt}^{H} \qquad \text{and} \qquad M_{kt}^* < M_{jt}^*.$$

B.4 Proof of Proposition 3

A firm in quality group k that resets its price in period t chooses a common reset price P_{kt}^*. If the firm is unable to reset again for ℓ consecutive periods, the price P_{kt}^* remains in effect in period $t+\ell$. The probability of this event is $(1-\zeta_k)^\ell$.

The demand of household type h for the firm's product in period $t+\ell$ is

$$D_{k,t+\ell}^{h}(P_{kt}^*) = \alpha_k (C_{t+\ell}^{h})^{\varepsilon_k} \left(\frac{P_{kt}^*}{E_{t+\ell}^{h}} \right)^{-\theta_{t+\ell}^{h}}.$$

Under the constant-discount-factor approximation, the firm's reset-price problem is

$$\max_{P_{kt}^*} \quad \mathbb{E}_t \sum_{\ell=0}^{\infty} [\beta(1-\zeta_k)]^\ell \left(P_{kt}^* - W_{t+\ell} \right) \sum_{h \in \{H,S\}} \lambda_h D_{k,t+\ell}^{h}(P_{kt}^*). \tag{B.10}$$

More generally, β^ℓ in (B.10) can be replaced by the appropriate nominal discount factor $Q_{t,t+\ell}$.

Since

$$\frac{\partial D_{k,t+\ell}^{h}(P)}{\partial P} = -\theta_{t+\ell}^{h} \frac{D_{k,t+\ell}^{h}(P)}{P},$$

the first-order condition is

$$0 = \mathbb{E}_t \sum_{\ell=0}^{\infty} [\beta(1-\zeta_k)]^\ell \sum_{h \in \{H,S\}} \lambda_h D_{k,t+\ell}^{h}(P_{kt}^*) \left[1 - \theta_{t+\ell}^{h} \frac{P_{kt}^* - W_{t+\ell}}{P_{kt}^*} \right]. \tag{B.11}$$

Equivalently,

$$P_{kt}^* = \frac{\mathbb{E}_t \sum_{\ell=0}^{\infty} [\beta(1-\zeta_k)]^\ell W_{t+\ell} \sum_h \lambda_h \theta_{t+\ell}^{h} D_{k,t+\ell}^{h}(P_{kt}^*)}{\mathbb{E}_t \sum_{\ell=0}^{\infty} [\beta(1-\zeta_k)]^\ell \sum_h \lambda_h (\theta_{t+\ell}^{h} - 1) D_{k,t+\ell}^{h}(P_{kt}^*)}.$$

Let μ^*_{kt} denote the log deviation of the flexible-price markup from its steady-state value. Log-linearising (B.11) around a zero-inflation steady state gives

$$p^*_{kt} = [1-\beta(1-\zeta_k)]\sum_{\ell=0}^{\infty}[\beta(1-\zeta_k)]^{\ell}\mathbb{E}_t\left(w_{t+\ell}+\mu^*_{k,t+\ell}\right).$$

Equivalently, the reset price satisfies the recursive equation

$$p^*_{kt} = [1-\beta(1-\zeta_k)]\left(w_t+\mu^*_{kt}\right)+\beta(1-\zeta_k)\mathbb{E}_t p^*_{k,t+1}. \tag{B.12}$$

The log-linearised price-index equation for quality group k is

$$p_{kt} = \zeta_k p^*_{kt} + (1-\zeta_k)p_{k,t-1}.$$

It follows that

$$\pi_{kt} = \zeta_k\left(p^*_{kt}-p_{k,t-1}\right). \tag{B.13}$$

Using (B.13) in (B.12) and rearranging yields

$$\pi_{kt} = \kappa_k\left(\mu^*_{kt}+w_{kt}\right)+\beta\mathbb{E}_t\pi_{k,t+1},$$

where

$$w_{kt}\equiv w_t - p_{kt}, \qquad \kappa_k = \frac{\zeta_k[1-\beta(1-\zeta_k)]}{1-\zeta_k}.$$

This proves Proposition 3.

B.5 Proof of Proposition 4

Log-linearising the first-order condition (B.4) around the steady state yields the following equation,

$$\alpha^{\mu^*}_k\mu^*_{kt}+\alpha^H_k(e^H_t-w_t)+\alpha^S_k(e^S_t-w_t)=0,$$

where

$$\alpha^{\mu^*}_k = \sum_{h=H,S}\lambda_h\left[-\alpha_k(\overline{C}^h)^{\varepsilon_k}\overline{\theta}^h(1-\overline{\theta}^h)(\overline{M}^*_k)^{-\overline{\theta}^h}(\overline{E}^h/\overline{W})^{\overline{\theta}^h}-\alpha_k(\overline{C}^h)^{\varepsilon_k}\overline{\theta}^h(1+\overline{\theta}^h)(\overline{M}^*_k)^{-\overline{\theta}^h-1}(\overline{E}^h/\overline{W})^{\overline{\theta}^h}\right], \tag{B.14}$$

$$\alpha^h_k = \lambda_h\left[(1-\overline{\theta}^h)\frac{\overline{C}^h_k}{\overline{E}^h}\frac{d\log C^h_k}{d\log E^h}-\delta_1\frac{\overline{\theta}^h\overline{C}^h_k}{\overline{E}^h}\frac{d\log C^h}{d\log E^h}+\overline{\theta}^h\frac{\overline{W}}{\overline{P}_k}\frac{\overline{C}^h_k}{\overline{E}^h}\frac{d\log C^h_k}{d\log E^h}+\delta_1\frac{\overline{W}}{\overline{P}_k}\frac{\overline{\theta}^h\overline{C}^h_k}{\overline{E}^h}\frac{d\log C^h}{d\log E^h}\right]\overline{E}^h.$$

Rearranging and simplifying these expressions yields equations (40), (41) and (B.17) in Proposition 4. The

elasticity of real consumption with respect to nominal expenditure, evaluated at the steady state, is

$$\frac{\partial \log C^h}{\partial \log E^h} = \frac{1-\overline{\theta}^h}{\overline{\varepsilon}^h + \delta_1 \overline{\theta}^h (\overline{e}^h - \overline{p}^h)} > 0, \tag{B.15}$$

where we define the steady-state expenditure-weighted average income-elasticity parameter and log price for household type h as

$$\overline{\varepsilon}^h \equiv \sum_{k=1}^{K} \overline{m}_k^h \varepsilon_k, \qquad \overline{p}^h \equiv \sum_{k=1}^{K} \int_{k-1}^{k} \overline{m}_{k\iota}^h \overline{p}_{k\iota}\, d\iota.$$

We next establish the sign of $\eta_k^{\mu^*}$. From the steady-state demand function (22), evaluated at the ideal price $\overline{P}_k = \overline{WM}_k^*$, we have

$$\overline{C}_k^h = \alpha_k (\overline{C}^h)^{\varepsilon_k} (\overline{M}_k^*)^{-\overline{\theta}^h} \left(\frac{\overline{E}^h}{\overline{W}}\right)^{\overline{\theta}^h}. \tag{B.16}$$

Substituting (B.16) into (B.14) gives

$$\alpha_k^{\mu^*} = \sum_{h\in\{H,S\}} \lambda_h \overline{\theta}^h \left[\left(\overline{\theta}^h - 1\right) - \frac{1+\overline{\theta}^h}{\overline{M}_k^*}\right] \overline{C}_k^h.$$

Therefore, defining $\eta_k^{\mu^*} \equiv -\alpha_k^{\mu^*}$, we obtain

$$\eta_k^{\mu^*} = \sum_{h\in\{H,S\}} \lambda_h \overline{\theta}^h \left[\frac{1+\overline{\theta}^h}{\overline{M}_k^*} - \left(\overline{\theta}^h - 1\right)\right] \overline{C}_k^h. \tag{B.17}$$

To establish positivity, define

$$X_k \equiv \lambda_H \overline{C}_k^H, \qquad Y_k \equiv \lambda_S \overline{C}_k^S, \qquad a \equiv \overline{\theta}^H, \qquad b \equiv \overline{\theta}^S.$$

Because $X_k > 0$ and $Y_k > 0$, the steady-state pricing condition (37) implies

$$\overline{M}_k^* = \frac{aX_k + bY_k}{(a-1)X_k + (b-1)Y_k}. \tag{B.18}$$

Substituting (B.18) into (B.17) and collecting terms yields

$$\eta_k^{\mu^*} = \frac{a(a-1)X_k^2 + q(a,b)X_kY_k + b(b-1)Y_k^2}{aX_k + bY_k}, \tag{B.19}$$

where

$$q(a,b) \equiv -a^2 + (4b-1)a - b^2 - b.$$

Under the maintained assumptions, $1 < b < a < 2b$. For a fixed $b > 1$, $q(a, b)$ is concave in a. Hence, over $a \in [b, 2b]$, its minimum is attained at one of the endpoints. At those endpoints, $q(b, b) = 2b(b-1) > 0$, and $q(2b, b) = 3b(b-1) > 0$.

It follows that $q(a, b) > 0$ whenever $b < a < 2b$. Since $a > 1$, $b > 1$, $X_k > 0$, and $Y_k > 0$, every term in the numerator of (B.19) is strictly positive, while its denominator is also strictly positive. Therefore, $\eta_k^{\mu^*} > 0$.

B.6 Additional Results in Section 5

B.6.1 Households' Expenditure Shares

We establish the following expenditure-share ordering in the zero-inflation steady state of the simplified Calvo model in Section 5. As in that section, there are no idiosyncratic productivity shocks, productivity is normalised to one, and all firms face the common nominal marginal cost $\overline{W}$. Assume that every quality group has a strictly positive probability of price adjustment.

In this steady state, firms within each quality group charge the same price. Since aggregate conditions are constant, a firm's optimal reset price coincides with its ideal price. Consequently,

$$\overline{P}_k = \overline{W}\,\overline{M}_k^*, \qquad \overline{P}_k^H = \overline{P}_k^S = \overline{P}_k.$$

The equality of household-specific sectoral price indices follows from the absence of within-group price dispersion.

Corollary B.1. *Consider the zero-inflation steady state described above, with $K \geq 2$ and $\alpha_k > 0$ for every quality group. Suppose that $\overline{C}^H < \overline{C}^S$, $\delta_1 < 0$, $\overline{\theta}^h > 1$ for $h \in \{H, S\}$, and the sufficient uniqueness condition in Proposition 1 holds for every group.*

If $\varepsilon_k < \varepsilon_j$, then

$$\frac{\overline{m}_k^H}{\overline{m}_j^H} > \frac{\overline{m}_k^S}{\overline{m}_j^S}. \tag{B.20}$$

If the lowest-quality group $\underline{k}$ and the highest-quality group $\overline{k}$ have, respectively, the strictly smallest and strictly largest values of ε_k, then

$$\overline{m}_{\underline{k}}^H > \overline{m}_{\underline{k}}^S, \qquad \overline{m}_{\overline{k}}^H < \overline{m}_{\overline{k}}^S. \tag{B.21}$$

Proof. Because there is no within-group price dispersion, household h's expenditure share on quality group k is

$$\overline{m}_k^h = \frac{\overline{P}_k\,\overline{C}_k^h}{\overline{E}^h} = \alpha_k (\overline{C}^h)^{\varepsilon_k} \left(\frac{\overline{P}_k}{\overline{E}^h}\right)^{1-\overline{\theta}^h}. \tag{B.22}$$

Taking the ratio across quality groups and then across household types gives

$$\frac{\overline{m}_k^H/\overline{m}_j^H}{\overline{m}_k^S/\overline{m}_j^S} = \left(\frac{\overline{C}^H}{\overline{C}^S}\right)^{\varepsilon_k-\varepsilon_j}\left(\frac{\overline{P}_k}{\overline{P}_j}\right)^{\overline{\theta}^S-\overline{\theta}^H}. \tag{B.23}$$

Suppose that $\varepsilon_k < \varepsilon_j$. Proposition 2 implies $\overline{M}_k^* < \overline{M}_j^*$. Since actual prices equal ideal prices in this steady state and marginal cost is common across groups, $\overline{P}_k < \overline{P}_j$. Moreover, $\overline{C}^H < \overline{C}^S$ and $\delta_1 < 0$ imply $\overline{\theta}^H > \overline{\theta}^S$. Hence,

$$\left(\frac{\overline{C}^H}{\overline{C}^S}\right)^{\varepsilon_k-\varepsilon_j} > 1, \qquad \left(\frac{\overline{P}_k}{\overline{P}_j}\right)^{\overline{\theta}^S-\overline{\theta}^H} > 1.$$

Equation (B.23) therefore establishes the relative expenditure-share ordering.

To establish the inequalities for the extreme quality groups, define $r_k \equiv \overline{m}_k^H/\overline{m}_k^S$. The relative expenditure-share ordering implies $r_k > r_j$ whenever $\varepsilon_k < \varepsilon_j$. Since expenditure shares are strictly positive and sum to one,

$$\sum_{k=1}^K \overline{m}_k^S r_k = \sum_{k=1}^K \overline{m}_k^H = 1, \qquad \sum_{k=1}^K \overline{m}_k^S = 1. \tag{B.24}$$

Thus, one is a positively weighted average of the ratios $\{r_k\}$. Under the strict ordering of the extreme groups, $r_{\underline{k}}$ is strictly larger than every other ratio and $r_{\overline{k}}$ is strictly smaller. It follows that

$$r_{\underline{k}} > 1 > r_{\overline{k}}, \tag{B.25}$$

which establishes $\overline{m}_{\underline{k}}^H > \overline{m}_{\underline{k}}^S$ and $\overline{m}_{\overline{k}}^H < \overline{m}_{\overline{k}}^S$. □

The corollary states that, in this steady state, hand-to-mouth households allocate a larger expenditure share to the lowest-quality group and a smaller share to the highest-quality group than savers. In the special case of two quality groups, with $\varepsilon_1 < \varepsilon_2$, it implies

$$\overline{m}_1^H > \overline{m}_1^S, \qquad \overline{m}_2^H < \overline{m}_2^S. \tag{B.26}$$

The zero-inflation steady-state restriction is important for this argument. Along a transition with staggered price adjustment, actual prices need not coincide with ideal prices, and within-group price dispersion can make sectoral price indices household-specific. Proposition 2 alone therefore does not establish the same expenditure-share ordering along an arbitrary transition. The corresponding expenditure patterns in the full quantitative model are evaluated numerically.

B.6.2 Discussion of Kimball Preferences

Kimball preferences generalise constant-elasticity-of-substitution (CES) preferences by allowing demand elasticity to vary with a variety's relative consumption. For a finite set of N goods, the consumption aggregator C is implicitly defined by

$$\frac{1}{N}\sum_{i=1}^{N}\Gamma\left(\frac{C_i}{C}\right)=1, \tag{B.27}$$

where

$$\Gamma(1)=1, \qquad \Gamma'(x)>0, \qquad \Gamma''(x)<0. \tag{B.28}$$

The factor $1/N$ normalises the measure of goods to one. At a symmetric allocation with $C_i = \overline{C}$ for all i, equation (B.27) implies $C = \overline{C}$, so that $C_i/C = 1$.

There are two distinctions between this demand system and the non-homothetic preferences in our model.

First, the Kimball aggregator is homogeneous of degree one and therefore represents homothetic preferences. At a fixed vector of goods prices, a proportional increase in expenditure scales all quantities and aggregate consumption proportionally. Consequently, relative quantities $x_i \equiv C_i/C$ and expenditure shares remain unchanged. If households share the same Kimball preferences and face the same prices, differences in expenditure alone do not generate differences in their relative consumption baskets.

Second, variable demand elasticity under Kimball preferences arises from the curvature of Γ. Under the standard monopolistic-competition assumption that an individual firm takes aggregate demand terms as given, its demand elasticity is

$$\theta(x)=-\frac{\Gamma'(x)}{x\Gamma''(x)}, \qquad x=\frac{C_i}{C}.$$

The restrictions $\Gamma'(x)>0$ and $\Gamma''(x)<0$ ensure that this elasticity is positive, but do not by themselves imply that it decreases with relative quantity. For example, the CES specification $\Gamma(x) = x^{(\overline{\theta}-1)/\overline{\theta}}$, with $\overline{\theta} > 1$, satisfies these restrictions and implies $\theta(x) = \overline{\theta}$.

A commonly used specification, as in Klenow and Willis (2016), instead implies

$$\theta(x)=\overline{\theta}x^{-\overline{\varepsilon}/\overline{\theta}}, \qquad \overline{\theta}>1, \quad \overline{\varepsilon}>0.$$

Here, $\overline{\theta}$ is the demand elasticity at $x = 1$, and $\overline{\varepsilon}$ governs its variation with relative quantity. In this specification,

$$\frac{d\log\theta(x)}{d\log x}=-\frac{\overline{\varepsilon}}{\overline{\theta}}<0.$$

Thus, holding aggregate demand terms fixed, a higher relative price reduces relative quantity and raises demand elasticity. Where $\theta(x) > 1$, this lowers the firm's desired markup. The CES case is recovered as $\overline{\varepsilon} \to 0$. The declining-elasticity property therefore follows from this additional functional-form restriction, rather than from monotonicity and concavity of Γ alone.

The distinction relevant to our mechanism is that homothetic Kimball preferences do not generate a direct effect of expenditure on demand elasticity at a fixed price vector: relative quantities, and hence $\theta(x_i)$, remain unchanged. This does not imply that Kimball preferences generate constant markups over the business cycle, because relative prices can change.

In our model, by contrast, household h's demand elasticity depends directly on its real aggregate consumption:

$$\theta_t^h = \exp\left(\delta_0 + \delta_1 \log C_t^h\right), \qquad \delta_1 < 0.$$

Even at a fixed price vector, a decline in household expenditure reduces real consumption and raises demand elasticity. Differences in expenditure across households therefore generate differences in both consumption baskets and price sensitivity. Changes in household expenditures affect firms' ideal markups through both household-specific elasticities and the composition of their customer base, as characterised in Section 5.3.

B.6.3 An Illustrative Example

To provide intuition for the compositional effect and the markup effect in Section 5.3, we provide an illustrative example. Consider an *inferior* good k with negative income-elasticity ($\psi_k < 0$). In Figure B.1, the light-grey areas represent the sales share to H consumers (denoted by ω^H), who exhibit a demand elasticity of θ^H. The dark-grey areas represent the sales share to S consumers (denoted by ω^S), whose demand elasticity is θ^S with $\theta^S < \theta^H$. If E^H declines while E^S remains unchanged, hand-to-mouth households allocate a greater share of their expenditure to good k. This increases ω^H (compositional effect) while simultaneously raising their price elasticity θ^H (markup effect). Together, these forces drive firm k to reduce its price, as shown in equation (41).

Figure B.1: Compositional and Markup Effect: An Example

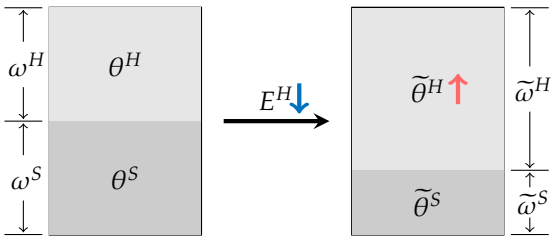

C. Appendix for Empirical Analysis

C.1 Data

C.1.1 Construction of Variables

This section details how we construct the variables used in the paper, including the dispersion of price changes, unit prices, price indices, and inflation.

Dispersion of price changes.
Let $\mathcal{O}_{c,t}$ be the set of retained non-zero UPC-store-week price-change observations in product module c assigned to month t. For observation j, with UPC i, store s, and week w, define $dp_{j,c,t} = \log P^R_{i,s,c,w} - \log P^R_{i,s,c,w-1}$, where P^R denotes the weekly regular price. Month assignment uses the date corresponding to the start of the recorded Stata week, implemented as `mofd(dofw(date_week))`. A UPC-store combination can contribute multiple weekly observations within a month. The interquartile range (IQR) of these pooled weekly changes is:

$$IQR_{c,t} = Q_3(dp_{j,c,t}) - Q_1(dp_{j,c,t}), \quad j \in \mathcal{O}_{c,t},$$

where $Q_3(\cdot)$ and $Q_1(\cdot)$ are the third and first quartiles of the retained weekly observations in the module-month. Let $N_{c,t} = |\mathcal{O}_{c,t}|$ count these observations, rather than distinct UPCs. The sample standard deviation (Std) is:

$$\sigma_{c,t} = \sqrt{\frac{\sum_{j\in\mathcal{O}_{c,t}}(dp_{j,c,t} - \overline{dp}_{c,t})^2}{N_{c,t}-1}},$$

where $\overline{dp}_{c,t} = N_{c,t}^{-1}\sum_{j\in\mathcal{O}_{c,t}} dp_{j,c,t}$ is the arithmetic mean of the retained weekly price changes within product module c and month t.

Finally, we aggregate $IQR_{c,t}$ and $\sigma_{c,t}$ across product modules to IQR_t and σ_t using revenue weights of product modules $\omega_{c,t}$:

$$\begin{aligned} IQR_t &= \sum_c \omega_{c,t} IQR_{c,t}, \\ \sigma_t &= \sum_c \omega_{c,t}\sigma_{c,t}, \end{aligned} \tag{C.1}$$

The dispersion at other disaggregated levels is defined similarly.

The unit price.
Because UPCs within a product module are extremely homogeneous in the dataset, it is reasonable to construct a quality index for each UPC within a product module in a store. To construct the quality variable for UPC i, we first calculate the unit price for i in store s in month t: $UP_{i,s,t} = P_{i,s,t}/prmult_i/multi_i/size_i$,

where $prmult_i$ is the price multiplier, which is used to indicate deals such as 3 for \$1, $multi_i$ shows how many of those goods appear in a given pack, and $size_i$ is the numeric size of the good.

C.1.2 Sales Filter

The sales filter is a variation of the algorithm in Midrigan (2011). Its purpose is to remove temporary price changes and recover a sequence of regular prices. We apply the filter separately to each UPC–store price trajectory. When multiple price observations occur for the same UPC–store in a given week, we retain the maximum price.

The algorithm has two parameters. We set the half-window length to $l = 5$, giving a nominal window of $2l + 1 = 11$ calendar weeks, and set the modal-frequency cutoff to $c = 0.33$. The filter is applied to annual input panels augmented with available observations from adjacent years.

For a given price trajectory within an input panel, let $\mathcal{T}$ denote the set of observed calendar weeks. For each $t \in \mathcal{T}$, define

$$\mathcal{W}_t = \{s \in \mathcal{T} : t - l \leq s \leq t + l\}. \tag{C.2}$$

Thus, the window contains the available observations within five calendar weeks on either side of the current week. Missing weeks do not cause the window to expand to include more distant observations.

The candidate regular price is the modal posted price in this window. When several prices share the highest frequency, we select the largest:

$$P_t^M = \max\left(\text{Mode}\left\{P_s : s \in \mathcal{W}_t\right\}\right). \tag{C.3}$$

We divide the number of occurrences of this modal price by the fixed nominal window length:

$$f_t = \frac{1}{2l+1} \sum_{s \in \mathcal{W}_t} \mathbb{1}\left\{P_s = P_t^M\right\} = \frac{1}{11} \sum_{s \in \mathcal{W}_t} \mathbb{1}\left\{P_s = P_t^M\right\}. \tag{C.4}$$

The denominator remains 11 when observations are unavailable near the boundaries of a price trajectory or input panel, or when some weeks are missing. We do not renormalise by the number of observed prices in the window. Consequently, the modal price must appear at least four times to exceed the cutoff.

The regular-price series is constructed recursively over observed weeks. Let $t_0 = \min \mathcal{T}$ be the first observed week in the trajectory within the input panel. We initialise the series using its modal price, regardless of the cutoff:

$$P_{t_0}^R = P_{t_0}^M. \tag{C.5}$$

For each subsequent observed week t, let

$$t^- = \max\{s \in \mathcal{T} : s < t\} \tag{C.6}$$

denote the preceding observed week. We then set

$$P_t^R = \begin{cases} P_t^M, & \text{if } f_t > c, \\ P_{t^-}^R, & \text{if } f_t < c. \end{cases} \tag{C.7}$$

Since f_t is an integer multiple of $1/11$, equality with $c = 0.33$ cannot occur. When the cutoff is not exceeded, the regular price is carried forward from the preceding observed week, including when intervening weeks are unobserved.

C.1.3 Regular price changes and measurement errors

To obtain robust results on price changes, we need to address the issue of time aggregation in the NielsenIQ data. Stores in the NRS data report to NielsenIQ weekly prices based on a Sunday to Saturday cycle, calculated as a ratio of revenue over quantity. This procedure introduces measurement error in price changes, because the actual pricing cycle in a store might be different from the Sunday-Saturday cycle. For example, suppose the price of an item was changed on Wednesday in a store. Two units of this product are sold in the same week, one on Monday with price p_1 and one on Friday with price p_2. The reported price to NielsenIQ is $(p_1 + p_2)/2$ (which is also what we get in the data set). Even though there is only one price change, we will obtain two consecutive smaller price changes, which will increase the frequency and decrease the size of price changes.

To address this issue, we apply a sales filter described above to posted prices to obtain regular prices. This procedure eliminates temporary price changes, including consecutive small price changes. Two consecutive small price changes due to time aggregation are replaced by one regular price change. The size and frequency of regular price changes are thus less affected.

C.1.4 Constructing Time-varying Dispersion of Price Changes

To track time-varying dispersion of price changes, we use two standard measures: the interquartile range (IQR) and the standard deviation (Std) of price changes.[1] Given the comprehensive coverage of the NRS dataset, multiple aggregation methods are possible. Following Vavra (2014), we first compute the IQR and Std of retained weekly regular-price changes pooled within each product category and month. We then aggregate these dispersion measures across product categories using sales-weighted market-level averages. Our findings remain robust to alternative aggregation methods, including at the sample level, store level, store-product-category level and product level.

[1]Following the literature, we focus on the distribution of non-zero regular price changes.

C.1.5 Variance Decomposition

Decomposition in Fact 1. Let j index a retained non-zero UPC-store-week observation, and let $dp_{j,s,c,t}$ denote its weekly log regular-price change in store s and category c, assigned to month t as described in Appendix C.1.1. Let $\omega_{j,s,c,t} \geq 0$ denote its observation weight, normalised such that

$$\sum_{s,c} \sum_{j \in O_{s,c,t}} \omega_{j,s,c,t} = 1$$

for every t, where $O_{s,c,t}$ is the set of retained weekly price-change observations pooled within store-category-month (s, c, t), and $N_{s,c,t} = |O_{s,c,t}|$. A UPC-store combination may contribute several observations in a month.

Define the store-category weight as

$$\omega_{s,c,t} \equiv \sum_{j \in O_{s,c,t}} \omega_{j,s,c,t},$$

and define the conditional weight of observation j within cell (s, c) as

$$\omega_{j|s,c,t} \equiv \frac{\omega_{j,s,c,t}}{\omega_{s,c,t}} = \frac{1}{N_{s,c,t}}.$$

Thus observations receive equal weight within each store-category-month. The cell weight $\omega_{s,c,t}$ is its annual revenue weight, normalised over cells observed in month t, and $\omega_{j,s,c,t} = \omega_{s,c,t}/N_{s,c,t}$. The weighted mean weekly price change in store-category cell (s, c) is

$$\overline{dp}_{s,c,t} \equiv \sum_{j \in O_{s,c,t}} \omega_{j|s,c,t} dp_{j,s,c,t},$$

while the economy-wide weighted mean at date t is

$$\overline{dp}_t \equiv \sum_{s,c} \omega_{s,c,t} \overline{dp}_{s,c,t}.$$

For any observation-level variable $x_{j,s,c,t}$, define the weighted variance over the weekly observations pooled in month t as

$$\mathrm{Var}(x \mid t) \equiv \sum_{s,c} \sum_{j \in O_{s,c,t}} \omega_{j,s,c,t} \left(x_{j,s,c,t} - \overline{x}_t\right)^2.$$

Similarly, define the conditional within-cell variance as

$$\mathrm{Var}(dp \mid s, c, t) \equiv \sum_{j \in O_{s,c,t}} \omega_{j|s,c,t} \left(dp_{j,s,c,t} - \overline{dp}_{s,c,t}\right)^2.$$

The weighted law of total variance then gives

$$\mathrm{Var}(dp \mid t) = \underbrace{\sum_{s,c} \omega_{s,c,t} \left(\overline{dp}_{s,c,t} - \overline{dp}_{t}\right)^2}_{\text{Between-store-category}} + \underbrace{\sum_{s,c} \omega_{s,c,t} \,\mathrm{Var}(dp \mid s,c,t)}_{\text{Within-store-category}}. \tag{C.8}$$

Equivalently, defining

$$\mathrm{Var}_{\omega_{s,c,t}}\left(\overline{dp}_{s,c,t} \mid t\right) \equiv \sum_{s,c} \omega_{s,c,t} \left(\overline{dp}_{s,c,t} - \overline{dp}_{t}\right)^2,$$

equation (C.8) can be written compactly as

$$\mathrm{Var}(dp \mid t) = \mathrm{Var}_{\omega_{s,c,t}}\left(\overline{dp}_{s,c,t} \mid t\right) + \sum_{s,c} \omega_{s,c,t} \,\mathrm{Var}(dp \mid s,c,t).$$

Decomposition in Fact 2. We next decompose the within-store-category variance into variation between quality groups and variation within quality groups. Fix a store-category-month cell (s,c,t), and let $O_{Q,s,c,t}$ denote the subset of weekly observations whose UPCs are assigned to quality group Q. Define $N_{Q,s,c,t} = |O_{Q,s,c,t}|$; these are observation counts, not counts of distinct UPCs.

The weight of quality group Q within cell (s,c,t) is

$$\omega_{Q|s,c,t} \equiv \sum_{j \in O_{Q,s,c,t}} \omega_{j|s,c,t} = \frac{N_{Q,s,c,t}}{N_{s,c,t}},$$

and the conditional weight of observation j within a non-empty quality group Q is

$$\omega_{j|Q,s,c,t} \equiv \frac{\omega_{j|s,c,t}}{\omega_{Q|s,c,t}} = \frac{1}{N_{Q,s,c,t}}.$$

The weighted mean price change of quality group Q is

$$\overline{dp}_{Q,s,c,t} \equiv \sum_{j \in O_{Q,s,c,t}} \omega_{j|Q,s,c,t} dp_{j,s,c,t}.$$

Define the weighted variance of quality-group mean price changes within cell (s,c,t) as

$$\mathrm{Var}_{\omega_{Q|s,c,t}}\left(\overline{dp}_{Q,s,c,t} \mid s,c,t\right) \equiv \sum_{Q} \omega_{Q|s,c,t} \left(\overline{dp}_{Q,s,c,t} - \overline{dp}_{s,c,t}\right)^2.$$

The weighted variance within quality group Q is

$$\mathrm{Var}(dp \mid s,c,Q,t) \equiv \sum_{j \in O_{Q,s,c,t}} \omega_{j|Q,s,c,t} \left(dp_{j,s,c,t} - \overline{dp}_{Q,s,c,t}\right)^2.$$

Applying the weighted law of total variance within each store-category-month cell gives

$$\operatorname{Var}(dp \mid s,c,t) = \underbrace{\operatorname{Var}_{\omega_{Q|s,c,t}}\left(\overline{dp}_{Q,s,c,t} \mid s,c,t\right)}_{\text{Between-quality-group}} + \underbrace{\sum_{Q} \omega_{Q|s,c,t} \operatorname{Var}(dp \mid s,c,Q,t)}_{\text{Within-quality-group}}. \tag{C.9}$$

C.1.6 Constructing Households' Recession Quality Index

To quantify household purchasing behaviour, we first construct a time-invariant market-level quality index for each product in the Minneapolis–St. Paul market. Let i denote a UPC–version combination, c its product category, and s a store. Let $Q_{i,s,c}$ denote the standardised store-level quality index defined in (1). We aggregate these store-level measures using fixed weights constructed from scanner revenues over 2006–2015.

Specifically, let $R_{i,s,c,y}$ denote annual revenue from product i at store s in year y, and let R_y denote total revenue across all products and stores in the market's retained scanner sample in that year. Define

$$a_{i,s,c,y} = \frac{R_{i,s,c,y}}{R_y}, \qquad \bar{\omega}_{i,s,c} = \frac{\sum_{y=2006}^{2015} a_{i,s,c,y}}{\sum_{r}\sum_{y=2006}^{2015} a_{i,r,c,y}},$$

where r indexes stores selling product i. Thus, the store weights are constructed by summing annual market-revenue shares and normalising the resulting sums across stores for each product.

The market-level quality index is

$$Q_{i,c} = \sum_{s \in \mathcal{S}_{i,c}^{Q}} \bar{\omega}_{i,s,c} Q_{i,s,c},$$

where $\mathcal{S}_{i,c}^{Q}$ comprises the stores whose observations for product i are matched between the revenue-weight and quality data. Both the store-level quality measures and the aggregation weights are fixed across quarters. Store weights are normalised before this match and are not renormalised after unmatched store–product observations are excluded.

We then construct a quarterly household quality index by weighting these fixed product scores by household expenditure shares. Let $\mathcal{M}_{h,c,t}$ denote the set of products in category c purchased by household h in quarter t that are matched to the market-level quality data. Let $e_{h,i,c,t}$ denote expenditure on product i, net of coupons. Define

$$\omega_{i,c,t}^{h} = \frac{e_{h,i,c,t}}{\sum_{c' \in C}\sum_{j \in \mathcal{M}_{h,c',t}} e_{h,j,c',t}},$$

where C is the set of product categories. These shares are normalised over the household's matched purchases across all categories, so that

$$\sum_{c \in C} \sum_{i \in \mathcal{M}_{h,c,t}} \omega^h_{i,c,t} = 1.$$

The household's quarterly quality index is therefore

$$Q_{h,t} = \sum_{c \in C} \sum_{i \in \mathcal{M}_{h,c,t}} \omega^h_{i,c,t} Q_{i,c}.$$

Changes in $Q_{h,t}$ reflect changes in expenditure shares across products with fixed quality scores. Consequently, the index measures changes in the composition of household expenditure evaluated using a common quality benchmark, rather than changes in the measured quality of a given product over time.

C.1.7 Households' Recession Quality Index and Income

In this section, we show that households' recession quality index defined in equation (6) is positively correlated with households' income, which shows that high-income households are more likely to purchase high-quality goods during recessions. The left panel in Figure C.1 plots the binscatter between households' recession quality index and log total expenditure. The right panel plots the binscatter between households' recession quality index and log total expenditure per capita. A one-standard-deviation increase in the index is associated with 0.1244 log points higher household expenditure and 0.1536 log points higher household-size-adjusted expenditure (approximately 12.4% and 15.4%, respectively).[2]

We use total expenditure to measure households' income because the nominal household incomes in the NielsenIQ Consumer Panel are measured imprecisely: incomes are reported in discrete income ranges with a two-year lag. When computing households' income per capita, we first regress log total expenditure on dummies for each household size and a full set of household controls. We next obtain household total expenditure per capita by subtracting the estimated coefficient associated with the relevant household-size indicator(s). Our results are robust even if we use the income definition in the NielsenIQ Consumer Panel to measure households' income.

[2]The percentage figures use the approximation $100b$, where b is the estimated coefficient in the log-expenditure regression. The exact transformation, $100[\exp(b) - 1]$, gives 13.2% and 16.6%, respectively.

Figure C.1: Households' Recession Quality Index and Income

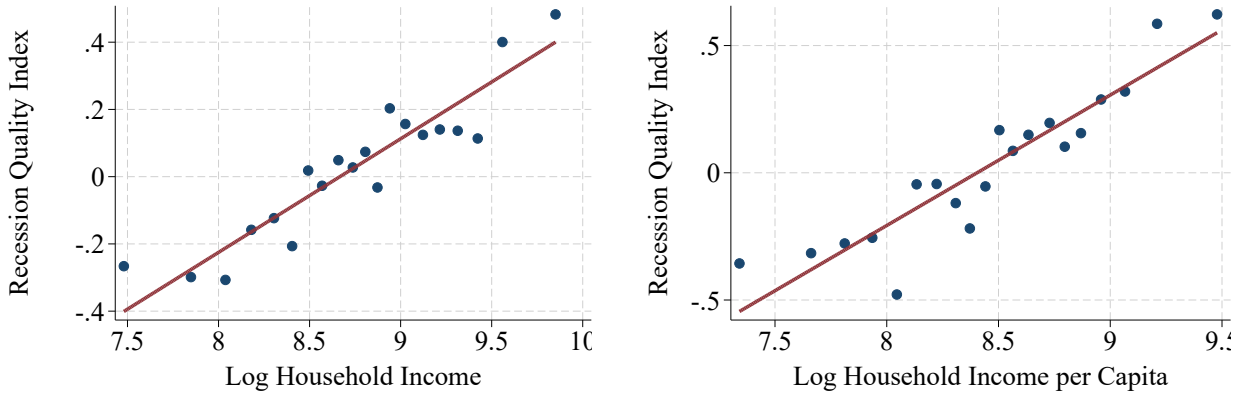

C.1.8 The Aggregate Frequency of Price Changes

Figure C.2 plots the aggregate frequency of regular price changes, the IQR of regular price changes, and the local unemployment rate. The aggregate frequency of price changes is countercyclical: the correlation between the frequency and the unemployment rate is 0.604. There is a significant positive correlation between the frequency and IQR of regular price changes, which is 0.387. These results are consistent with findings in Vavra (2014).

Figure C.2: The Aggregate Frequency of Price Changes

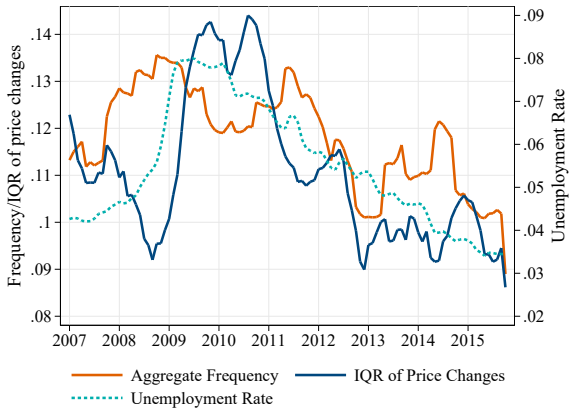

Notes: This figure plots the aggregate frequency of regular price changes, the IQR of regular price changes, and the local unemployment rate (seasonally adjusted) of Minneapolis-St. Paul MSA. The unemployment data are downloaded from the Bureau of Labor Statistics. The aggregate frequency is calculated using the NielsenIQ Retail Scanner data. Before seasonal adjustment and smoothing, monthly frequency is constructed as $f_m = 1 - \prod_{w \in \mathcal{W}_m}(1 - f_w)$, where f_w is the revenue-weighted aggregate weekly frequency of regular price changes and $\mathcal{W}_m$ contains the weeks assigned to month m using each week's start date. The IQR series is constructed from retained weekly regular-price changes pooled within each month. For better graphical illustration, we move IQR of price changes downwards by 0.08.

C.2 Summary Statistics of Price Changes

Table C.1 reports summary statistics of regular-price adjustment behaviour in our baseline sample. The empirical size and distribution statistics use retained weekly regular-price changes pooled over the sample; adjustment frequency is aggregated separately to months. Several features of the data are noteworthy. First, the average frequency of price adjustment is 12.9% per month, with a clear quality gradient: low-quality goods adjust more frequently (15.6%) than high-quality goods (8.7%). Second, the absolute size of price changes declines monotonically with quality, from 14.6% for low-quality goods to 11.9% for high-quality goods. A similar pattern emerges for the standard deviation of price changes.

Third, the distribution of price changes exhibits substantial dispersion and fat tails. The kurtosis exceeds three for the full sample and for most quality groups, indicating deviations from normality. At the same time, a large fraction of adjustments are small: roughly 41% of price changes are smaller than one-half of the mean absolute price change. The percentile statistics further illustrate the skewed and dispersed nature of price adjustments, with a median absolute change of 9% and an interquartile range spanning from 3% to 18% in the full sample. Overall, price adjustment behaviour varies systematically with product quality and displays significant heterogeneity in both frequency and size.

Table C.1: Summary Statistics of Regular Price Changes

	Full sample	By Quality				
Statistics		Low	Medium-low	Medium	Medium-high	High
Frequency	0.129	0.156	0.140	0.123	0.112	0.087
Absolute size	0.135	0.146	0.140	0.133	0.123	0.119
Standard deviation	0.200	0.219	0.205	0.195	0.183	0.177
Kurtosis	3.05	3.14	2.92	2.96	3.15	3.11
Fraction small	0.41	0.42	0.41	0.40	0.41	0.42
Absolute size distribution						
25th percentile size	0.03	0.03	0.04	0.04	0.03	0.03
50th percentile size	0.09	0.10	0.10	0.09	0.08	0.08
75th percentile size	0.18	0.20	0.19	0.18	0.17	0.16

Notes: Size and distribution statistics are calculated from retained weekly regular-price changes pooled over the sample. Within each quality group, frequency is computed by summing revenue-weighted weekly adjustment indicators within each month and averaging the monthly totals. The five quality groups use unweighted pooled quantile cutoffs of unit-price ranks normalised by the maximum rank within each store-category cell, using average unit prices over 2006–2015. Group sizes need not be equal within a cell or in the matched sample. "Fraction small" represents the fraction of small price changes, where small price changes are defined as those smaller than (1/2)mean($|dp|$).

C.3 Countercyclical Dispersion of Price Changes

Figure C.3 documents a salient cyclical property of price adjustments. The figure plots monthly interquartile range (IQR) measures constructed from pooled weekly regular-price changes together with the unemploy-

ment rate in the Minneapolis–St. Paul MSA. A clear countercyclical pattern emerges: the dispersion of price changes rises during periods of elevated unemployment and declines during expansions.

Figure C.3: Countercyclical Dispersion of Price Changes

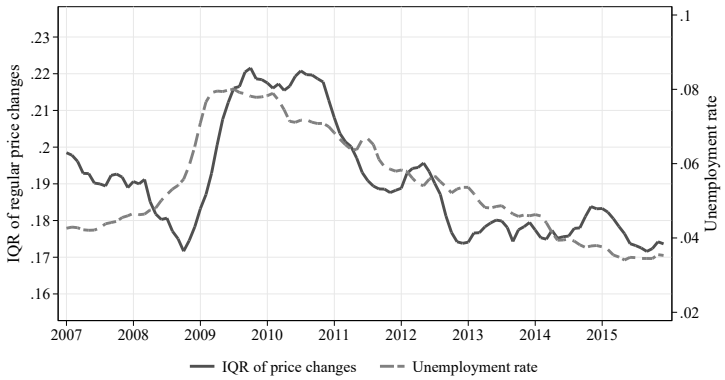

Notes: This figure plots monthly interquartile range (IQR) measures of retained weekly regular-price changes pooled within each month alongside the unemployment rate in the Minneapolis-St. Paul MSA. The unemployment rate is sourced from the Bureau of Labor Statistics. All variables are seasonally adjusted.

C.4 Frequency and Size of Price Changes: Decomposition

Following Klenow and Kryvtsov (2008), we examine the positive and negative frequency and magnitude of price changes for the lowest and highest of four quality groups. We rank UPCs by their average unit price over 2006–2015 within each store-category cell, divide by the maximum rank in the cell, and apply unweighted pooled quartile cutoffs to the normalised ranks. These cutoffs need not produce equal-sized groups within each cell. Figure C.4 plots groups 1 and 4; the two middle groups are not shown. This is a separate supporting exercise, rather than an exact decomposition of the five-group inflation regressions in the main text. The difference in the cyclicality of inflation between high- and low-quality goods is primarily driven by the frequency of price changes (extensive margin): in the Great Recession, for high-quality goods, the frequency of positive price changes drops slightly, while the frequency of negative price changes remains roughly constant. For low-quality goods, the frequency of positive price changes drops, and the frequency of negative price changes increases dramatically. The average size of price changes, conditional on adjustment, varies little over time for both quality groups.

Figure C.4: Frequency and Size Decomposition (Quality)

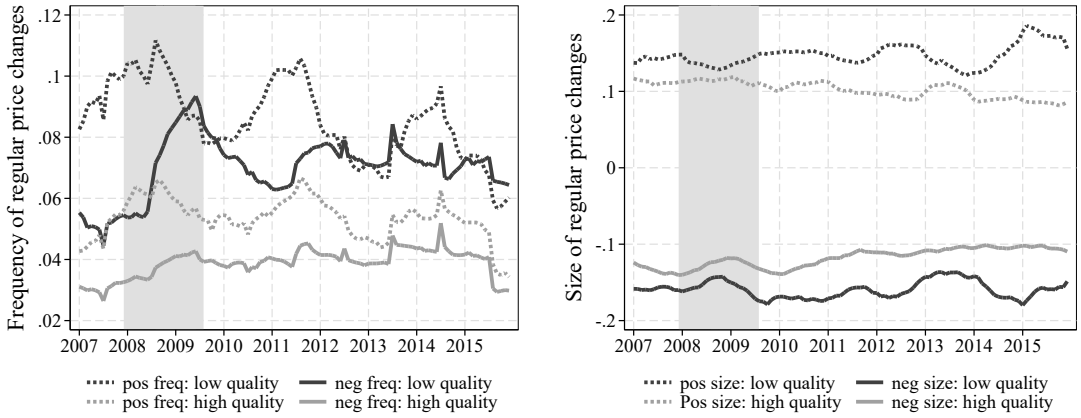

(a) Frequency: High and Low Quality

(b) Size: High and Low Quality

Notes: This figure shows the frequency and size of price changes for the lowest and highest of four quality groups. The groups use pooled quartile cutoffs of normalised within-store-category unit-price ranks over 2006–2015; groups 2 and 3 are not shown. These groups differ from the five groups in the main-text inflation regressions. Panel (a) decomposes the frequency of price changes to the frequency of positive price changes and the frequency of negative price changes. Panel (b) decomposes the size of price changes to its positive and negative components.

C.5 Inflation Dynamics during 2017–2023

Our baseline analysis focuses on the Great Recession period (2006–2015). We now extend the analysis to cover 2017–2023, which includes the COVID-19 recession and the subsequent high-inflation episode.

Figure C.5 plots the annual inflation rates for the lowest- (solid blue) and highest-quality (dashed red) groups. The shaded area marks the early-2020 recession, when unemployment spiked and real GDP growth fell sharply (Figure C.6, panel (a)). To cushion the downturn, the U.S. government enacted large fiscal stimulus packages, indicated by the vertical dashed lines in Figure C.5: the CARES Act ($2 trillion) and the American Rescue Plan ($1.9 trillion). These amounts far exceeded the roughly $900 billion stimulus during the Great Recession. The unprecedented transfers boosted aggregate demand and helped the economy rebound quickly in 2021, but they also contributed significantly to the post-2021 inflation surge. Recent studies suggest that fiscal stimulus accounted for a substantial share of the demand-driven inflation in this period. Di Giovanni et al. (2023) decompose U.S. inflation during 2019–2022 and find that aggregate-demand shocks explain about two-thirds of the rise, with fiscal policy responsible for roughly half or more of that demand component. Similarly, Bernanke and Blanchard (2025) argue that, absent the large fiscal and monetary support that created exceptionally strong demand, supply-side disruptions alone would not have generated sustained inflation after 2021.

Why, then, did low-quality goods experience much higher inflation than high-quality goods during the post-2021 surge, as seen in Figure C.5? Our model offers a mechanism: low-income households received transfers that represented a larger share of their income and became more price inelastic compared to high-income households. Low-income households also have higher marginal propensities to consume, and

spend most of their stimulus payments, with a disproportionate share flowing towards low-quality goods. As a result, retailers raise markups more for those goods than for high-quality goods, leading to higher inflation for low-quality items.[3]

Finally, note that while aggregate inflation dipped slightly during the 2020 recession, food inflation actually rose (Figure C.6, panel (b)). Our results are consistent with this fact: because food accounts for a large share of sales in the NielsenIQ data, we observe slightly positive overall inflation in 2020. As Figure C.5 shows, low-quality goods also experienced marginally higher inflation than high-quality goods during 2020, though the difference was relatively small.

Figure C.5: Annual Inflation: Lowest and Highest Quality

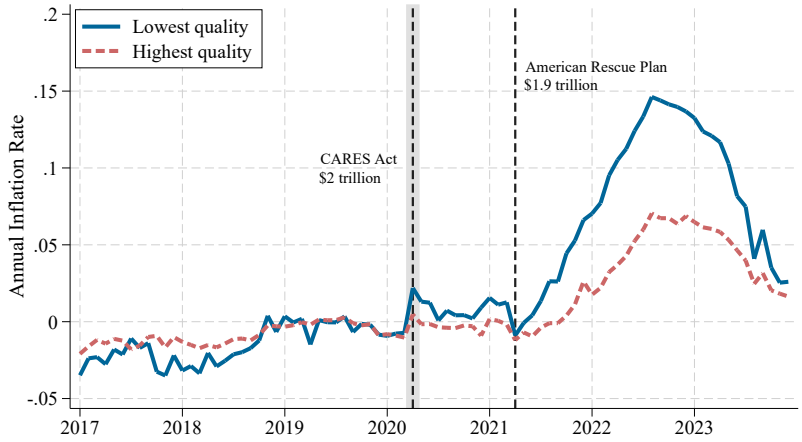

Notes: This figure plots annual inflation from 2017 to 2023 for the lowest and highest of five quality groups. These groups use unweighted pooled quintile cutoffs of unit-price ranks normalised by the maximum rank within each store-category cell, with average unit prices calculated over 2016–2020. Group assignments remain fixed over the plotted period and need not yield equal group sizes within each cell. All inflation rates are computed using annual UPC-store revenue weights. The annual inflation is defined as $\pi_{k,t} = p_{k,t} - p_{k,t-12}$, where t denotes a month. The shaded area denotes the 2020 Recession. Vertical dashed lines indicate the CARES Act (late March/April 2020) and the American Rescue Plan (March 2021).

[3]A recent paper by Mongey and Waugh (2025) presents a related argument, showing quantitatively that fiscal stimulus can lead to divergent price adjustments across quality groups.

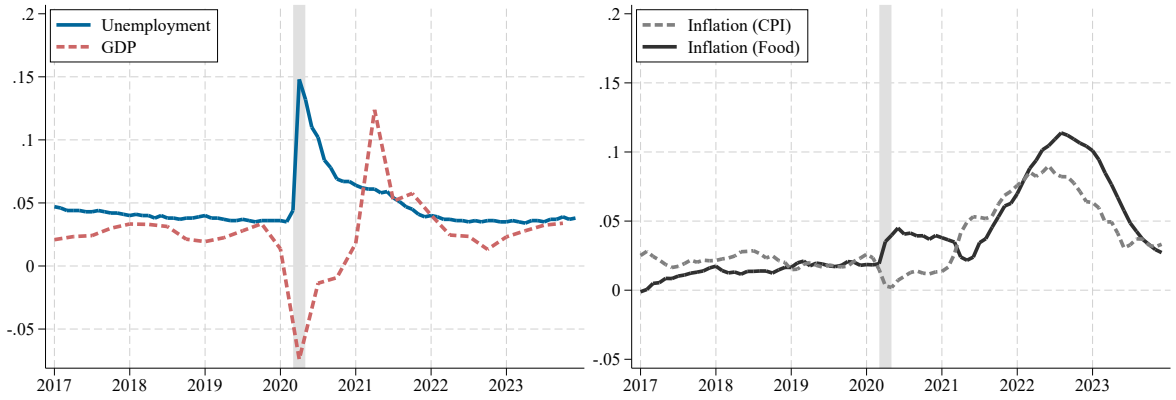

(a) GDP and Unemployment (b) Inflation

Figure C.6: Aggregate Conditions during the COVID-19 Recession

Notes: Panel (a) plots U.S. real GDP growth and the unemployment rate. Panel (b) plots overall CPI inflation and food inflation. Shaded areas denote the 2020 recession.

C.6 Empirical Results in Atlanta-Sandy Springs MSA

For manageability, we conduct our baseline analysis in the Minneapolis-St. Paul MSA. In this section, we examine the baseline patterns in the Atlanta-Sandy Springs MSA using a restricted store sample. Stores are ranked within each MSA and retailer parent by cumulative revenue over 2006–2015. The sample construction retains up to the three highest-revenue stores in each MSA–parent group and then applies a 20% random retention probability to these stores. The replication uses the retained store list supplied with the processed data, without drawing a new sample. The Atlanta results therefore describe this restricted sample rather than all stores in the market.

Countercyclical Dispersion of Price Changes. Figure C.7 is the counterpart of Figure C.3 mentioned in Section 3. Following the same monthly pooling of retained weekly regular-price changes, it plots the interquartile range (IQR) of those observations (black solid line) and the unemployment rate (grey dashed line) in the Atlanta market. It clearly shows that the dispersion of price changes is countercyclical — it is strongly positively correlated with the local unemployment rate.

Figure C.7: Countercyclical Dispersion of Price Changes

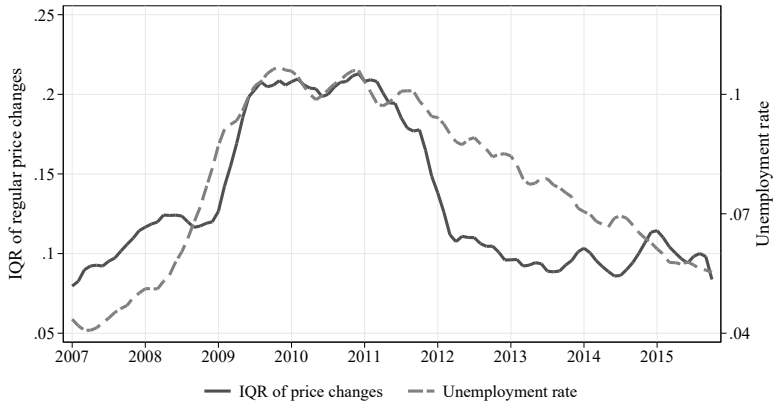

Notes: This figure plots monthly interquartile range (IQR) measures of retained weekly regular-price changes pooled within each month alongside the unemployment rate in the Atlanta-Sandy Springs MSA. The unemployment rate is sourced from the Bureau of Labor Statistics. All variables are seasonally adjusted.

Variance decomposition. Figure C.8 applies the variance decomposition of Figure 1 in Section 3 to the common quality-matched Atlanta sample. Panel (a) shows that both within-store-category and between-store-category variance increase during the Great Recession. Panel (b) further decomposes the same within-store-category component and shows that between-quality variance contributes to its increase, alongside within-quality variance. The increase in between-quality dispersion provides qualitative support for the baseline finding in this restricted sample.

Figure C.8: Variance Decomposition

(a) Variance between and within store-category

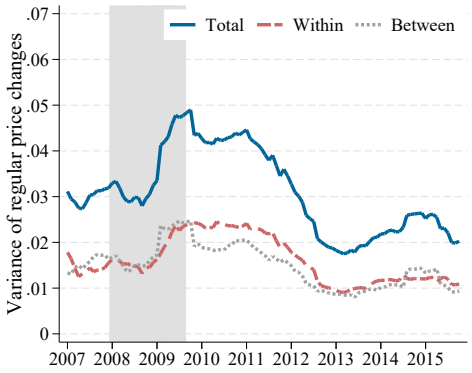

(b) Variance between and within quality

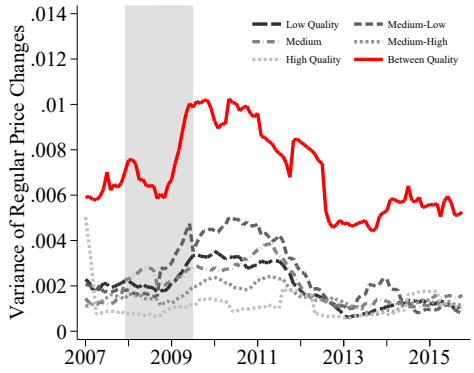

Notes: Both panels use the same Atlanta sample of weekly regular-price changes with a valid store–UPC quality assignment and an annual store-category revenue weight. Weekly observations are pooled within each month, and a UPC-store combination may contribute multiple observations in a month. Store selection is described above. Panel (a) decomposes total variance into within-store-category and between-store-category components using equation (2). Panel (b) decomposes exactly the within-store-category component in panel (a) into between-quality variance and the five within-quality contributions. Its five quality groups use unweighted pooled quintile cutoffs of normalised within-store-category unit-price ranks constructed from the Atlanta 2006–2015 data. The cutoffs are specific to Atlanta and do not impose equal group sizes within each cell. Observations receive equal weight within each store-category-month, so quality-group weights equal their shares of those observations. Both panels use the same annual store-category revenue weights, renormalised over observed cells in each month. All series use a nine-month centred moving average and are displayed from January 2007 through October 2015.

C.7 Robustness

C.7.1 Correlation with GDP

In our baseline, we illustrate the countercyclical dispersion of price changes by showing that it is positively correlated with the local unemployment rate. We choose to use the local unemployment rate mainly because it is collected monthly but the local GDP data are annual. Here, we check the robustness of our findings using GDP. In Figure C.9, we plot annual averages of the adjusted and smoothed monthly dispersion series alongside annual log(GDP), which are clearly negatively correlated. The correlation is -0.23.

Figure C.9: Countercyclical Dispersion of Price Changes

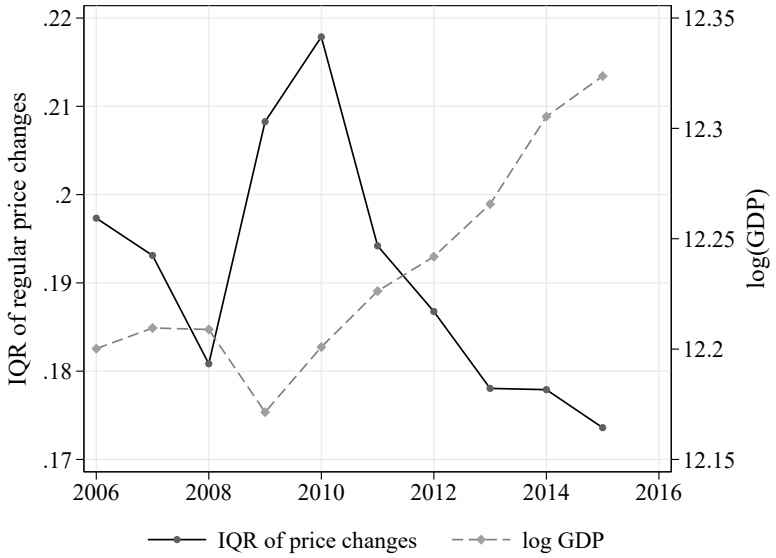

Notes: This figure plots annual averages of the seasonally adjusted and smoothed monthly interquartile range (IQR) series alongside log(GDP) for the Minneapolis-St. Paul MSA. Each monthly IQR is constructed from retained weekly regular-price changes pooled within that month and aggregated across product modules using revenue weights. The GDP data are downloaded from U.S. Bureau of Economic Analysis.

D. Appendix for Quantitative Analysis

D.1 Numerical Solution Methods

In Section 6, we numerically solve for the transition path of the perfect foresight equilibrium in response to unexpected demand shocks and monetary policy shocks. We first briefly outline the solution method. We assume that the economy starts in the steady state and returns to it after 200 periods. We then conjecture the entire path of household inflation $\{\pi_t^S, \pi_t^H\}$ and wage inflation over the transition path. Given these conjectured prices, we solve the households' problem backward and derive firms' pricing policy functions. Next, we simulate the population of firms forward using a non-stochastic simulation algorithm similar to Young (2010), generating the distribution over individual state variables at each period. Using this distribution, we compute the sequence of prices and inflation rates. If these computed prices and inflation rates differ from the conjectured ones, we update the guessed prices iteratively until the equilibrium converges.

Initial guess. We assume that the economy returns to its steady state after $T = 200$ periods. Let $\overline{\pi}^h$ and $\overline{\pi}^W$ denote the steady-state household and wage inflation rates, respectively. We initialise the transition paths as

$$\pi_t^{h,(0)} = \overline{\pi}^h, \qquad \pi_t^{W,(0)} = \overline{\pi}^W \tag{D.1}$$

The converged equilibrium is invariant to these initial guesses.

Solving households' problem backward. For a given set of inflation $\{\pi_t^S, \pi_t^H, \pi_t^w\}$, we solve households' problem backward in the following steps:

1. calculate the aggregate inflation π_t.[4]
2. obtain the path of nominal interest rate i_t and real interest rate r_t from the Taylor rule and the Fisher equation.
3. solve for the consumption path $\{C_t^S, C_t^H\}$ and labour supply $\{N_t^S, N_t^H\}$ using savers' Euler equation and savers' and hand-to-mouth households' intratemporal optimality condition and budget constraints.
4. calculate households' elasticity of substitution $\{\theta_t^S, \theta_t^H\}$ once we have the consumption path.

Solving firms' problem backward. Given the set of inflation $\{\pi_t^S, \pi_t^H, \pi_t^w\}$ and households' consumption, we use value-function iteration to solve firms' price-setting problem for their policy functions.[5] We also use the MacQueen-Porteus Error Bounds to improve speed. We use the method in Tauchen (1986) to discretise the AR(1) process.

[4] We normalise the nominal wage at period -1.

[5] Knotek and Terry (2008) compare grid-based methods such as value function iteration with collocation methods in state-dependent pricing models. They find grid-based methods to be more robust.

In calculating the steady-state equilibrium, whenever an equilibrium is solved, the algorithm checks the grid range to make sure it is neither too large so that many grid points are wasted nor too small so that the bounds of the grid points are reached. Den Haan (2010) reports that algorithms that use the largest range for individual capital also have lower accuracy. The lower and upper bounds of the grid are chosen in order not to miss or waste grid points. Otherwise, the algorithm dynamically adjusts the grids so that the former criterion is met. This step is crucial to obtain the accurate solution of the model as well as accurately simulated moments of price adjustment.

Simulating population of firms forward. We next simulate firms' pricing decisions using the policy functions from the previous step. We adopt the *non-stochastic simulation method* in Young (2010) to simulate the population of firms.[6] This simulation generates the distribution over firms' individual state variables at each period. Using this distribution, we can compute the sequence of aggregate prices.

Updating and Convergence. The simulation above gives the updated household-specific inflation $\{\widetilde{\pi}_t^S, \widetilde{\pi}_t^H\}$. We use the following steps to obtain the updated wage inflation $\{\pi_t^w\}$,

1. calculate the aggregate inflation π_t.[7]
2. obtain the path of nominal interest rate i_t and real interest rate r_t from the Taylor rule and the Fisher equation.
3. solve for the consumption path $\{C_t^S, C_t^H\}$, labour supply $\{N_t^S, N_t^H\}$ and nominal wage W_t jointly using savers' Euler equation and savers' and hand-to-mouth households' intratemporal optimality condition and budget constraints. In particular, we can express the total nominal profits in the economy as a function of C_t^S and C_t^H.
4. obtain updated $\{\pi_t^w\}$ from the wage sequence $\{W_t\}$.

If these updated inflation sequences differ from the ones in the previous step, we update the guessed prices iteratively until the equilibrium converges.

D.2 Steady-State Model Fit

Table D.1 evaluates the model's ability to match key steady-state moments across quality groups. The table reports model-implied moments alongside their empirical counterparts. Moments are organised into three categories: (a) targeted expenditure moments, (b) targeted pricing moments, and (c) non-targeted pricing moments. The empirical pricing moments use the five-group store-level classification described in the notes to Table C.1. The expenditure groups in panel (a) are constructed separately: UPCs are ordered by market-level unit price within each product category, assigned the integer percentile $1 + \lfloor 100(j-1)/N_c \rfloor$, where j is the position in that ordering and N_c is the number of UPCs in the category, and then allocated

[6] Algan et al. (2008) show that there are two potential drawbacks of stochastic simulation. First, there is sampling noise. Second, the realisations of the random draw cluster around the mean, leading to large approximation error at the tail.

[7] We normalise the nominal wage at period -1.

to five groups using unweighted pooled quintile cutoffs of these percentiles. Both classifications have five ordered groups, but their product membership need not coincide.

Panel (a) shows that the model closely replicates the cross-sectional expenditure shares of hand-to-mouth (H) and saver (S) households across quality groups. The model matches both the level and gradient of expenditure shares, capturing the concentration of hand-to-mouth spending in lower-quality goods and the relatively more even distribution of saver expenditures. These moments discipline the non-homothetic demand structure and pin down the heterogeneity in price elasticity across goods.

Panel (b) reports the pricing moments targeted in the calibration: frequency of price adjustment, average absolute size, kurtosis, and the fraction of small price changes. The empirical size and distribution moments in panels (b) and (c) use retained weekly regular-price changes pooled over the sample; empirical frequency is aggregated to months as described in Table C.1. The model reproduces the declining frequency of price changes with quality and closely matches the magnitude of price adjustments across groups. It also captures the non-Gaussian features of price changes, including the broad kurtosis pattern and the substantial share of small adjustments. These moments jointly identify the parameters governing nominal rigidity and the distribution of idiosyncratic shocks.

Panel (c) presents moments not directly targeted in the estimation. The model performs well along these dimensions, matching the standard deviation and the cross-sectional distribution of absolute price changes across quality groups. In particular, the model captures the broad decline in dispersion as quality increases. The ability to replicate these non-targeted moments provides additional validation of the joint discipline imposed by the demand system and the pricing structure. Overall, the steady-state fit indicates that the model successfully captures both the expenditure heterogeneity and the rich cross-sectional patterns in price adjustment observed in the data.

Table D.1: Model Fit: Steady State

Moments	Quality Group									
	Model					Data				
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
(a) Expenditure moments targeted										
Expenditure share (H)	0.28	0.33	0.20	0.13	0.06	0.28	0.33	0.20	0.13	0.06
Expenditure share (S)	0.19	0.32	0.22	0.17	0.10	0.19	0.32	0.22	0.17	0.10
(b) Pricing moments targeted										
Frequency	0.16	0.14	0.12	0.12	0.09	0.16	0.14	0.12	0.11	0.09
Absolute size	0.16	0.14	0.13	0.14	0.11	0.15	0.14	0.13	0.12	0.12
Kurtosis	3.05	2.83	3.34	3.17	3.94	3.14	2.92	2.96	3.15	3.11
Fraction small	0.36	0.38	0.40	0.39	0.45	0.42	0.41	0.40	0.41	0.42
(c) Moments not targeted (price changes)										
Std. dev. of price changes	0.21	0.18	0.18	0.18	0.15	0.22	0.21	0.19	0.18	0.18
Absolute size distribution										
25th percentile size	0.03	0.02	0.02	0.02	0.01	0.03	0.04	0.04	0.03	0.03
50th percentile size	0.14	0.13	0.11	0.12	0.08	0.10	0.10	0.09	0.08	0.08
75th percentile size	0.25	0.22	0.20	0.22	0.17	0.20	0.19	0.18	0.17	0.16

D.3 The Role of Non-homotheticity

D.3.1 Results with Homothetic Preferences

We construct an alternative model in which we replace non-homothetic preferences in our baseline model with homothetic CES preferences ($\delta_1 = 0$; $\varepsilon_k = 1 - \exp(\delta_0)$), so that all households share the same constant elasticity of substitution $\exp(\delta_0)$ and consume identical baskets of goods. This economy is subjected to the same demand shocks specified in Section 6.2.

Figure D.1: Homothetic CES: Dispersion of Price Changes

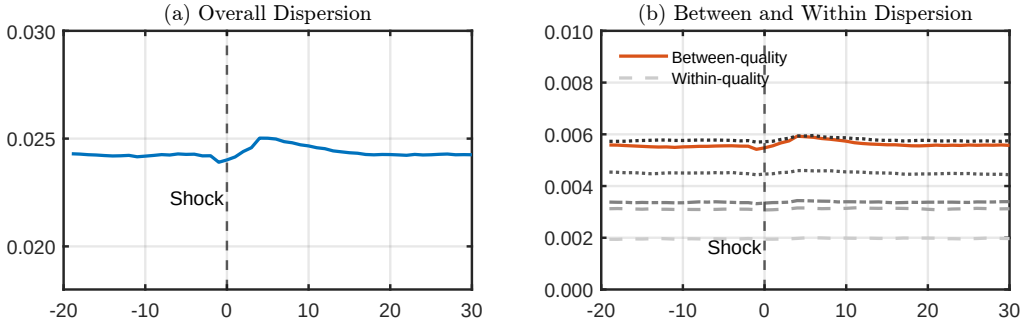

Notes: Horizontal axes denote months. Vertical axes denote price-change variance. The series are smoothed using a centred 5-month window. The overall dispersion, the between and within dispersion are averaged over 10 simulations.

Figure D.1 plots the total variance (panel (a)) as well as the between-group and within-group variance (panel (b)). Unlike our baseline model, the homothetic version fails to generate enough countercyclical dispersion in price changes. Figure D.2 further contrasts households' price elasticity in the homothetic model with the baseline model, and contrasts the inflation differential between quality groups in the homothetic model against their data counterparts. Since homothetic preferences imply a constant price elasticity over time, as shown in panel (a), inflation and its cyclicality across quality groups follow similar paths, as shown in panel (b) and panel (c), which are too small compared to the data. This therefore gives rise to negligible changes in the between-group dispersion following the demand shocks.

Figure D.2: Homothetic CES vs. Baseline: Price Elasticity and Inflation

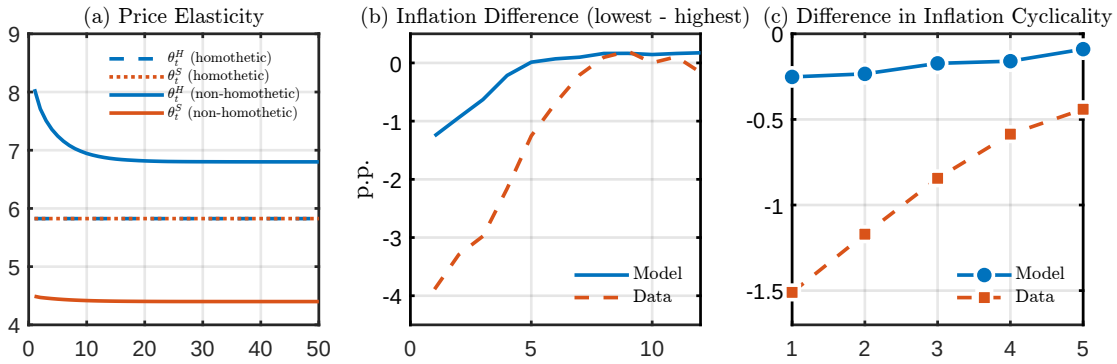

D.3.2 Channel Decomposition

The non-homothetic CES preference in our model deviates from the homothetic CES preference in two respects: non-homotheticity in income elasticity and non-homotheticity in price elasticity. Which channel is essential to replicate our empirical findings in price dispersion and inflation dynamics? To answer this question, we conduct the following experiment: we shut down the non-homotheticity in the income elasticity while retaining the non-homothetic price elasticity.

To shut down the channel of income elasticity and only allow for variations in price elasticity, we introduce a household-specific component $\nu_{k,t}^{h,\varepsilon}$ to the income-elasticity parameter $\varepsilon_{k,t}^h = \varepsilon_k + \nu_{k,t}^{h,\varepsilon}$.

In the context of our baseline model, the demand is now given by

$$C_{k\iota t}^h = \alpha_k (C_t^h)^{\varepsilon_{k,t}^h + \theta_t^h - 1} \left(\frac{P_{k\iota t}}{P_t^h} \right)^{-\theta_t^h} C_t^h,$$

where

$$P_t^h = \left(\sum_{k=1}^{K} \alpha_k (C_t^h)^{\varepsilon_{k,t}^h + \theta_t^h - 1} \int P_{k\iota t}^{1-\theta_t^h} d\iota \right)^{\frac{1}{1-\theta_t^h}}.$$

In each period for household h, we choose the value $\nu_{k,t}^{h,\varepsilon}$ such that $\varepsilon_{k,t}^h + \theta_t^h - 1 = 0$ so that the demand function

is identical to the one in homothetic CES preference except that price elasticity varies with consumption according to

$$\log \theta_t^h = \log \bar{\theta}^h + \delta_1 \log \left(\frac{C_t^h}{\bar{C}_{\text{PE}}^h} \right).$$

Here, $\bar{C}_{\text{PE}}^h$ denotes household h's steady-state consumption in this counterfactual. We retain each household's baseline steady-state price elasticity, $\bar{\theta}^H = 6.8$ and $\bar{\theta}^S = 4.4$, and the baseline slope δ_1, which governs the response of log price elasticity to proportional consumption changes. The elasticity rule is therefore centred on each household's counterfactual steady-state consumption. Equivalently, it has household-specific intercepts $\delta_{0,\text{PE}}^h = \log \bar{\theta}^h - \delta_1 \log \bar{C}_{\text{PE}}^h$.

Figure D.3: Price-Change Dispersion: Price Elasticity Channel

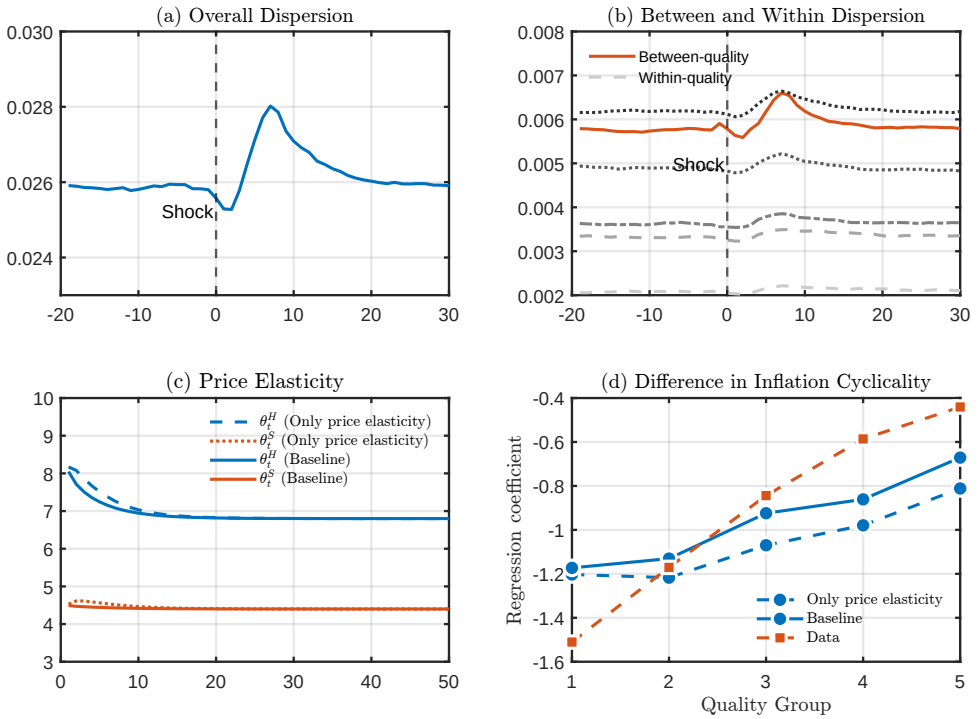

Notes: In panels (a) – (c), horizontal axes denote months. In panels (a) and (b), vertical axes denote price-change variance, and the series are smoothed using a centred 5-month window. The overall dispersion, the between and within dispersion are averaged over 10 simulations.

Panel (a) and panel (b) in Figure D.3 plot the overall as well as the between-group and within-group dispersion of price changes. The results change little compared to our baseline model. Panel (c) displays price elasticity responses similar to those in the baseline model. Panel (d) shows that the inflation cyclicality between quality groups is mainly driven by the channel of non-homothetic price elasticity, as the model predictions align closely with those in our baseline model.

D.4 The Role of Price Stickiness

We first examine a simplified flexible-price benchmark that retains two household types, five quality groups, and non-homothetic demand. In addition to removing nominal rigidity, it eliminates idiosyncratic firm productivity and within-group price dispersion, normalises productivity to unity, and has one freely adjusting price per quality group. Nominal marginal cost is therefore the common wage, W_t. The benchmark uses its own steady state and the shock and elasticity parameters specified below. It illustrates the distinction between markup levels and inflation dynamics; it does not isolate the quantitative effect of price stickiness with all other features held fixed. We then examine how the approximate group-level Phillips curves describe the baseline sticky-price model.

D.4.1 Inflation Dynamics under Flexible-Price Equilibrium

Demand shocks are multiplicative preference shifters, denoted here by ψ_t^h, with marginal utility $\psi_t^h(C_t^h)^{-\eta}$, as in the baseline. The pre-shock values are $\psi_0^h = 1$. On impact, $\psi_1^H = 0.780$ and $\psi_1^S = 0.961$; subsequently, $\psi_t^h = 1 + 0.8(\psi_{t-1}^h - 1)$. Price elasticities satisfy $\theta_t^h = \bar{\theta}^h(C_t^h/\bar{C}^h)^{-2.1}$, with $\bar{\theta}^H = 6.8$ and $\bar{\theta}^S = 4.4$. Bars denote this benchmark's steady state. These choices retain the baseline shock specification and calibration but replace its elasticity-response parameter; they are not a recalibration that holds all other features fixed.

Given the real allocation path, let $q_t^h = P_t^h/W_t$ denote household h's price index relative to the wage. Household inflation and the aggregate inflation measure used in monetary policy are calculated consistently as

$$1 + \pi_t^h = (1 + \pi_t^w)\frac{q_t^h}{q_{t-1}^h}, \qquad 1 + \pi_t = \prod_{h\in\{H,S\}} (1 + \pi_t^h)^{1/2}.$$

The nominal block is solved backward after incorporating the non-homothetic Euler wedge into transformed returns. Let $I_t = 1 + i_t$ denote the actual gross nominal return and $R_t = I_t/(1 + \pi_{t+1}^S)$ the actual gross real return measured in units of savers' consumption. Under perfect foresight, the saver Euler equation is

$$\psi_t^S(C_t^S)^{-\eta} = \beta R_t V_t \psi_{t+1}^S(C_{t+1}^S)^{-\eta}, \qquad \eta \equiv \frac{1}{\vartheta} = 1.$$

Here, V_t is the non-homothetic wedge defined in the main text, evaluated using this benchmark's preferences and allocation:

$$V_t = (1 + \pi_{t+1}^S)\frac{\partial C_{t+1}^S/\partial E_{t+1}^S}{\partial C_t^S/\partial E_t^S},$$

where C_t^S denotes savers' indirect real-consumption function, with derivatives evaluated holding the corresponding period's goods prices fixed.

We retain the Taylor-rule coefficients $\varphi_\pi = 1.4$, $\varphi_y = 0$, and $\varphi_v = -1$. With $V^* = 1$, inflation target $\pi^* = \bar{\pi}$,

and no monetary-policy shock, the rule for the actual nominal return is

$$I_t = \frac{1+\bar{\pi}}{\beta}\left(\frac{1+\pi_t}{1+\bar{\pi}}\right)^{1.4}\frac{1}{V_t},$$

where $\bar{\pi} = 0.002$ is monthly steady-state inflation. Define the transformed returns

$$\widetilde{R}_t \equiv R_t V_t, \qquad \widetilde{I}_t \equiv I_t V_t.$$

The Euler equation, Fisher relation, and Taylor rule then become

$$\widetilde{R}_t = \frac{\psi_t^S}{\beta\psi_{t+1}^S}\left(\frac{C_{t+1}^S}{C_t^S}\right)^{\eta}, \qquad \widetilde{I}_t = \widetilde{R}_t(1+\pi_{t+1}^S),$$
$$\widetilde{I}_t = \frac{1+\bar{\pi}}{\beta}\left(\frac{1+\pi_t}{1+\bar{\pi}}\right)^{1.4}.$$

Thus, the inflation-only expression applies to the transformed nominal return $\widetilde{I}_t$. The wedge cancels when the Euler equation and Taylor rule are combined. Actual returns can be recovered as

$$I_t = \frac{\widetilde{I}_t}{V_t}, \qquad R_t = \frac{\widetilde{R}_t}{V_t}.$$

The simulation has 200 transition periods, followed by a terminal return of consumption and relative prices to steady state and shock multipliers to one. Terminal wage growth is $1+\bar{\pi}$; terminal household inflation also incorporates the final relative-price adjustment. Household inflation is updated along the entire path, including the saver inflation entering the Fisher relation.

For the following two identities, inflation denotes log price changes, wage inflation denotes log wage growth, and $\mu_{kt}^* = \log\left(M_{kt}^*/\bar{M}_k^*\right)$. With unit productivity, the flexible price is $P_{kt} = W_t M_{kt}^*$, so

$$\pi_{kt}^{\text{flex}} = \mu_{kt}^* - \mu_{k,t-1}^* + \pi_t^w = \Delta\mu_{kt}^* + \pi_t^w. \tag{D.2}$$

Consequently, the inflation difference between any two groups j and k is

$$\pi_{jt}^{\text{flex}} - \pi_{kt}^{\text{flex}} = \Delta\mu_{jt}^* - \Delta\mu_{kt}^*. \tag{D.3}$$

Thus, differences in log inflation are entirely driven by differences in log markup growth. For the simple monthly rates plotted in Figure D.4, the exact accounting identity is $1+\pi_{kt}^{\text{simple}} = (1+\pi_t^{w,\text{simple}})M_{kt}^*/M_{k,t-1}^*$; adding simple markup growth and wage inflation is only a first-order approximation.

Figure D.4 displays the impulse responses following demand shocks in panel (a). Hand-to-mouth households experience a sharper cyclical decline in consumption (panel (b)), which makes them more price-elastic in the recession (panel (c)). Because these households concentrate their spending on low-quality goods, the

ideal markups of low-quality goods fall more sharply than those of high-quality goods (panel (d)).

Figure D.4: Impulse Responses to Demand Shocks

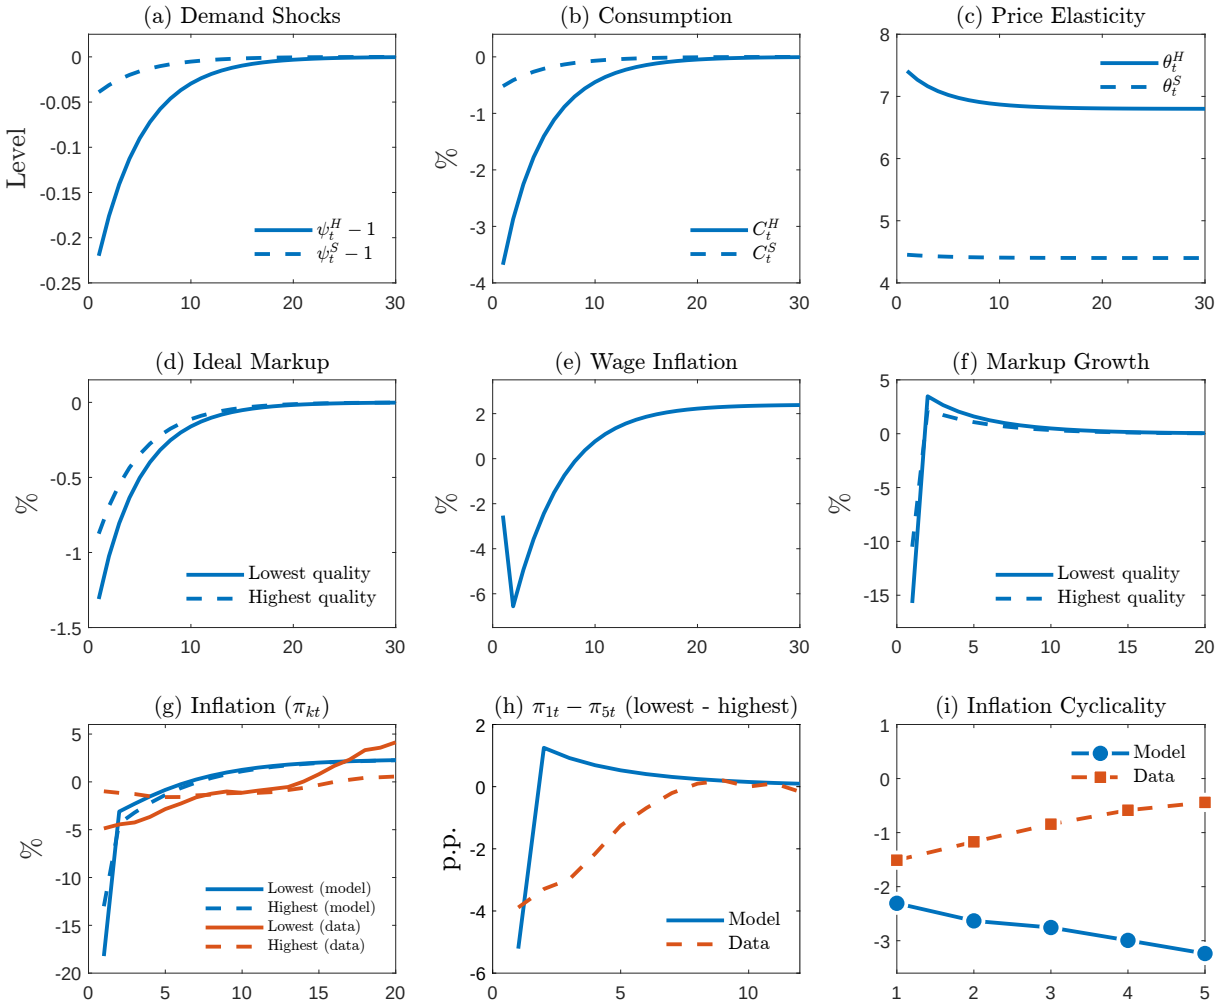

Notes: Horizontal axes denote months, except in panel (i), which denotes quality groups. Shocks hit in month 1. Panel (a) plots $\psi_t^h - 1$; panel (b) reports consumption deviations from steady state in percent; panel (c) reports elasticity levels. Panel (d) reports $100(M_{kt}^*/\bar{M}_k^* - 1)$, without annualisation. Model wage inflation, markup growth, and goods inflation in panels (e), (f), and (g) are simple monthly rates multiplied by 1200, expressed in annualised percent. Panel (h) reports the annualised inflation difference in percentage points. The data in panel (g) cover January 2009–August 2010; those in panel (h) cover January–December 2009. Annualisation by a factor of 12 is an approximation, not compounding.

Guided by equation (D.2), panels (e) and (f) plot wage inflation and ideal-markup growth, respectively, expressed as annualised monthly percentage changes. On impact, ideal markups fall more sharply for low-quality goods. After the initial period, ideal markups begin to mean-revert, and their growth rates are positive and larger for low-quality goods. Because wage inflation is common across quality groups, equation (D.3) implies that flexible-price inflation is also higher for low-quality goods during this recovery. Panel (g) displays annualised monthly model inflation for the lowest-quality group (solid lines) and highest-quality group (dashed lines), with the model predictions in blue and the empirical counterparts in red.

Unlike the data, the flexible-price model predicts higher inflation for low-quality goods than for high-quality goods. This qualitative reversal after the impact period is summarised in panel (h), which shows the inflation difference between the lowest- and highest-quality groups ($\pi_{1t} - \pi_{5t}$) over the twelve months following the shock. Panel (i) provides a descriptive comparison of the model's cyclicality diagnostic with the empirical estimates in Figure 2 of Section 3. These results are conditional on the simplified benchmark and its parameterisation.

D.4.2 The Approximate Phillips Curves

Within this simplified benchmark, ideal markups fall on impact and recover afterwards. Their recovery generates positive markup growth and reverses the ordering of inflation across quality groups after the initial period. This exercise illustrates how markup growth enters flexible-price inflation; because other model features also change, it does not establish that nominal rigidity alone accounts for the contrast with the data.

In contrast, the inflation dynamics in our baseline sticky-price model are well explained by the approximate Phillips curve in equation (39):

$$\pi_{kt} = \kappa_k(\mu^*_{kt} + w_{kt}) + \beta\, \mathbb{E}_t \pi_{k,t+1}.$$

Iterating forward gives

$$\pi_{kt} = \kappa_k \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t \left(\mu^*_{k,t+s} + w_{k,t+s}\right), \tag{D.4}$$

where the current inflation rate depends on the discounted sum of all current and *future* deviations of the ideal markup from steady state. Following the negative demand shocks, these deviations remain persistently negative (Figure D.5, panel (a)), thereby generating persistent deflation. Moreover, the ideal markup falls much more for low-quality goods, because hand-to-mouth households are more adversely affected by the recession.

Figure D.5: Inflation and the Approximated Phillips Curves

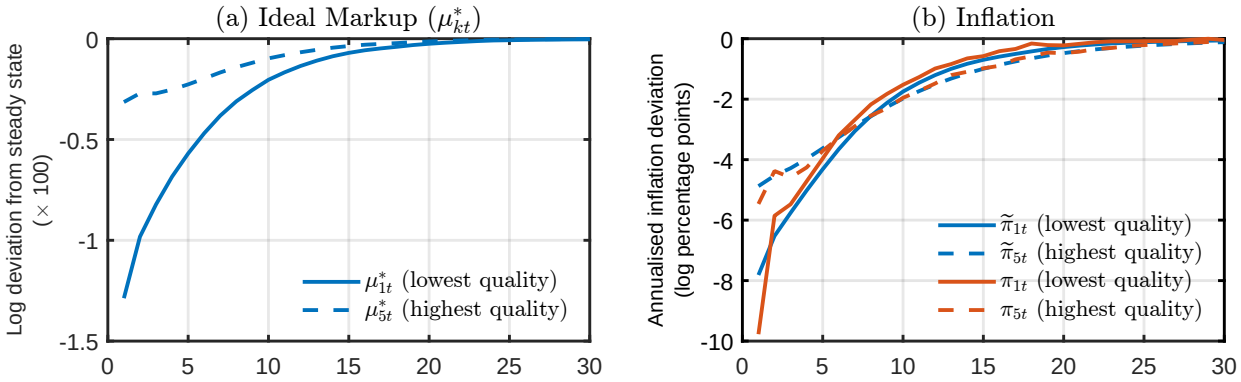

Notes: Horizontal axes denote months. Shocks hit in month 1. Panel (a) reports log deviations of ideal markups from their steady-state values, multiplied by 100. Panel (b) reports log inflation deviations from steady state, multiplied by 1200 (annualised log percentage points). Solid lines represent low-quality goods, while dashed lines represent high-quality goods. The variable $\widetilde{\pi}$ denotes inflation implied by the approximate Phillips curve, whereas π denotes inflation in the full nonlinear baseline model. Ideal markups are computed for a firm with productivity normalised to one, solving the markup fixed point along the saved baseline household consumption, expenditure, and elasticity paths. All deviations use the corresponding steady-state values from the same saved baseline solution. The approximate Phillips curve is solved backward over all 200 saved months with a zero terminal inflation deviation.

Using expression (D.4), we construct the quality-group inflation rates in the Calvo model, denoted by $\widetilde{\pi}_{kt}$. Panel (b) shows that the inflation rates π_{kt} from the full nonlinear baseline model (red lines) closely align with the linearised approximations $\widetilde{\pi}_{kt}$ from the approximate Phillips curves (blue lines).[8]

[8]Equation (D.4) is derived by log-linearising a model with Calvo pricing, whereas our baseline model is solved globally with fixed adjustment costs. Thus, while the dynamics are very similar, π_{kt} and $\widetilde{\pi}_{kt}$ do not coincide exactly. We choose the frequency of price changes in the Calvo model to match the IRFs of the baseline model.

E. Additional Tables and Figures

Table E.1: Monthly Inflation, Unemployment Rate and Quality of Goods (Quantity Weight)

	Monthly Inflation						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Quality (Std)×	0.426***	0.418***	0.178***	0.426***	0.155***	0.44	4.27***
Unemp Rate	(0.077)	(0.043)	(0.032)	(0.044)	(0.041)	(0.928)	(0.758)
Unemp Rate	-0.619***	-0.813***		-0.879***		-3.06***	0.067
	(0.057)	(0.033)		(0.034)		(0.902)	(0.509)
Quality (Std)×						-0.0006***	-0.0005***
Wholesale Inflation						(0.0001)	(0.0001)
Wholesale Inflation						-0.004***	-0.004***
						(0.0003)	(0.0003)
Fixed effects							
Store-Category		✓					
Store-Category-Time			✓				
Firm-Store-Category				✓			✓
Firm-Store-Category-Time					✓		

Notes: The dependent variable is the annualized monthly inflation. Retail inflation and unemployment are seasonally adjusted; the event-based wholesale cost-change measure in columns (6) and (7) is not. These columns use the full wholesale-control specification described in the main text; the quality main effect and the unemployment–wholesale cost-change interaction are included but their coefficients are not reported. Heteroskedasticity-robust standard errors are reported in columns (1) and (6); standard errors are clustered at the store-category level in columns (2)–(5) and (7). Column (7) includes firm–store–category fixed effects, whereas column (6) does not; the two specifications are otherwise identical. The specification in column (7) controls for persistent differences across firm–store–category units, but does not absorb time-varying firm-level cost shocks; the paper's inference about differences in inflation cyclicality across quality levels rests on column (7). All specifications include a constant that is not reported. *Significance*: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table E.2: Households' Recession Quality Index and Shopping Behaviour (Quality Index)

	Dependent Variable							
	Log(Spending)		Share Deal		Share Generic		Share Coupon	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Unemployment Rate	-0.015***	0.045	0.016***	0.019	0.007***	0.011*	0.003***	0.007
	(0.005)	(0.062)	(0.001)	(0.017)	(0.001)	(0.006)	(0.000)	(0.006)
$R_h^{Recession}\times$								
Unemployment Rate	0.032***	0.044***	-0.005**	-0.009***	-0.005***	-0.005***	-0.001	-0.001
	(0.009)	(0.009)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
Household Income		0.100***		0.000		0.005**		0.002
		(0.027)		(0.008)		(0.002)		(0.002)
Household Income ×								
Unemployment Rate		-0.008		0.000		-0.000		-0.001
		(0.005)		(0.001)		(0.000)		(0.001)
Household controls		✓		✓		✓		✓
Household controls×								
unemployment rate		✓		✓		✓		✓
Household FE	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.625	0.630	0.818	0.821	0.675	0.688	0.639	0.648
Observations	43339	43339	43339	43339	43339	43339	43339	43339

Notes: All specifications are the same as those in Section 3. The only difference is that the quality index is based on log relative unit prices rather than standardised unit prices. Standard errors are clustered at the household level. The regression is weighted by household projection weights provided by the NielsenIQ Consumer Panel. Even-numbered columns include household controls and interactions of all controls with the unemployment rate. Household controls comprise categorical indicators for the age and presence of children, the ages and education levels of the female and male household heads, marital status, and race, together with log household income adjusted for household size, which enters continuously. *Significance*: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table E.3: Households' Recession Quality Index and Shopping Behaviour (Dummy)

	Dependent Variable							
	Log(Spending)		Share Deal		Share Generic		Share Coupon	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Recession	0.053***	0.023	-0.002	0.027	0.001	0.009	0.001	-0.021
	(0.019)	(0.256)	(0.003)	(0.050)	(0.001)	(0.016)	(0.002)	(0.026)
$\mathbf{1}_{Low-Quality} \times \mathbf{1}_{Recession}$	-0.038	-0.028	0.009*	0.013**	0.002	0.003**	0.006**	0.004
	(0.025)	(0.027)	(0.005)	(0.005)	(0.001)	(0.001)	(0.002)	(0.002)
Household Income		0.077***		0.010**		0.003***		0.006***
		(0.018)		(0.004)		(0.001)		(0.002)
Household Income × Recession		-0.000		-0.005		-0.001		0.001
		(0.021)		(0.004)		(0.001)		(0.002)
Household controls		✓		✓		✓		✓
Household controls × $\mathbf{1}_{Recession}$		✓		✓		✓		✓
Household FE	✓	✓	✓	✓	✓	✓	✓	✓
R^2	0.502	0.511	0.813	0.816	0.676	0.680	0.595	0.618
Observations	34885	34885	34885	34885	34885	34885	34885	34885

Notes: The recession indicator equals one from 2007Q4 through 2009Q2. The low-quality indicator equals one for households whose recession quality index is strictly below the 33rd percentile and zero for households whose index is strictly above the 66th percentile. These cutoffs are calculated using one observation per household. Households in the middle group and those tied at either cutoff are excluded from the regression sample. Even-numbered columns include household controls and interactions of all controls with the recession indicator. Household controls comprise categorical indicators for the age and presence of children, the ages and education levels of the female and male household heads, marital status, and race, together with log household income adjusted for household size, which enters continuously. Household income and its interaction with the recession indicator are reported separately. All specifications include household fixed effects. Regressions are weighted using the projection weights provided by the NielsenIQ Consumer Panel. Standard errors are clustered at the household level. *Significance*: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure E.1: Impulse Responses (Baseline Model)

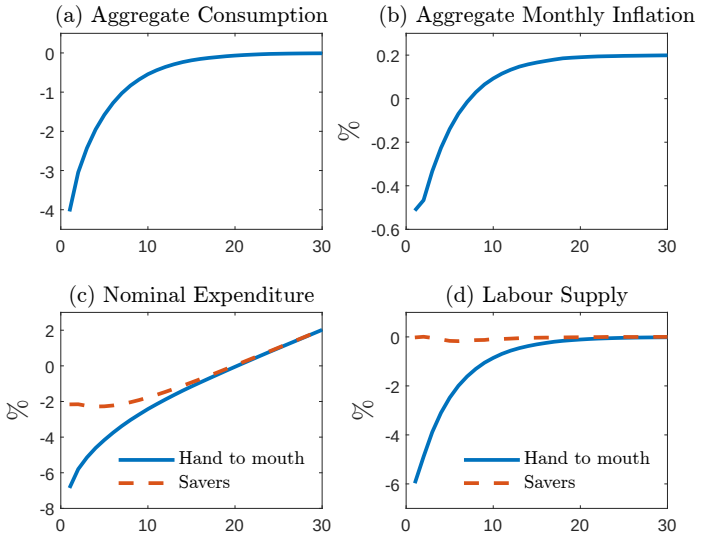

Notes: Horizontal axes denote months. Shocks hit in month 1. The vertical axes in panels (c) and (d) show percent deviations from steady state. Panel (c) reports percentage changes in nominal expenditure relative to its pre-shock level. Panel (d) reports percentage deviations in labour supply from steady state.

Supplementary References

Algan, Y., O. Allais, and W. J. Den Haan (2008): "Solving heterogeneous-agent models with parameterized cross-sectional distributions," *Journal of Economic Dynamics and Control*, 32, 875–908.

Bernanke, B., and O. Blanchard (2025): "What caused the US pandemic-era inflation?" *American Economic Journal: Macroeconomics*, 17, 1–35.

Den Haan, W. J. (2010): "Comparison of solutions to the incomplete markets model with aggregate uncertainty," *Journal of Economic Dynamics and Control*, 34, 4–27.

Di Giovanni, J., Ş. Kalemli-Özcan, A. Silva, and M. A. Yldrm (2023): "Quantifying the inflationary impact of fiscal stimulus under supply constraints," in *AEA Papers and Proceedings* Volume 113, 76–80, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Klenow, P. J., and O. Kryvtsov (2008): "State-dependent or time-dependent pricing: Does it matter for recent US inflation?" *The Quarterly Journal of Economics*, 123, 863–904.

Knotek, E. S., and S. Terry (2008): "Alternative methods of solving state-dependent pricing models," *Federal Reserve Bank of Kansas City Working Paper*.

Midrigan, V. (2011): "Menu costs, multiproduct firms, and aggregate fluctuations," *Econometrica*, 79, 1139–1180.

Tauchen, G. (1986): "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics letters*, 20, 177–181.

Vavra, J. (2014): "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *The Quarterly Journal of Economics*, 129, 215–258.

Young, E. R. (2010): "Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations," *Journal of Economic Dynamics and Control*, 34, 36–41.