Do Poor Households Pay Higher Markups in Recessions?

Jonathan Becker[†]

First draft: January 2024

This draft: February 2025

Abstract

Poor and rich households differ greatly in the mix of products they consume, with the poor allocating a larger share of their spending to relatively inexpensive goods. Moreover, during recessions, households shift spending toward more affordable goods. In this paper, I study an economy with nonhomothetic preferences and endogenously variable markups that is calibrated to match these patterns. I show that in recessions, producers of low-quality goods gain market power and increase markups because consumers shift spending toward more affordable goods. By contrast, producers of higher-quality goods reduce their markups. Observed changes in the expenditure distribution during the Great Recession predict a 6.8-percentage-point increase in the markups of low-quality goods and a 1.8-percentage-point decline in the markups of high-quality products, considerably increasing real consumption inequality. Embedding this mechanism into a Bewley-Aiyagari-Hugget model, I find redistributive policies aimed at alleviating inequality amplify these unequal markup movements. Redistribution to the poor allows lower-quality producers to gain even more market share and to increase markups even further.

Keywords: Markups, Inequality, Quality, Recessions, Nonhomotheticities.

JEL Codes: D11, D21, D43, E31, E32, L13.

[†] I owe a heartfelt debt of gratitude to my advisors Virgiliu Midrigan, Corina Boar, Raquel Fernández, and Mark Gertler for inspiration, continued guidance, and tireless support throughout this project. I thank Guido Menzio, Alessandra Peter, Jarda Borovička, Simon Gilchrist, Venky Venkateswaran, Diego Perez, Thomas Philippon, Paolo Martellini, Danial Lashkari, Ricardo Lagos, Tommy Iao, Abdou Ndiaye, Katka Borovičková, Fernando Cirelli, Gian Luca Clementi, Nobu Kiyotaki, Paco Buera, Chris Edmond, Benny Kleinman, Sergio Salgado, Paul Scott, and Jess Benhabib for helpful discussions and comments. Researcher(s)' own analyses 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 analyzing and preparing the results reported herein.
NYU; Email: jb7026@nyu.edu.

1 Introduction

Poor and rich households are known to systematically differ in their spending on cheap versus premium goods. Additionally, during recessions, all households shift expenditures toward more affordable goods.[1] Consequently, when a recession hits, producers of cheap versus premium goods face different changes in demand. This paper asks: How do markups respond to these unequal demand shifts in an environment in which producers' market power increases with their market share? And how does the corresponding change in relative prices affect households across the expenditure distribution?

I answer these questions by studying a model economy with two key ingredients: nonhomothetic preferences and oligopolistically competitive sectors. In each sector, households purchase a basket of imperfectly substitutable varieties. I interpret nonhomotheticities over these varieties as a quality margin, whereby the value households place on quality varies with their spending. Specifically, households with higher expenditures allocate a larger share of their spending to more expensive, higher-quality varieties and vice versa. I show that the cross-sectional heterogeneity in consumption patterns and price elasticities implied by these variety-level nonhomotheticities is closely aligned with the micro data.

Sectors are oligopolistically competitive, as in Atkeson and Burstein (2008). As a result, firms derive their market power from the imperfect substitutability of their products. While smaller firms primarily compete within their sectors, facing higher substitutability and lower market power, larger firms dominate their sectors and primarily compete across sectors. This cross-sector competition, where outputs are less substitutable, allows larger firms to wield greater market power and charge higher markups. As a result, markups are endogenously variable and increase with a firm's market share within its sector.

[1] Bils and Klenow (1998) show luxury spending is more cyclical. Burstein, Eichenbaum, and Rebelo (2005) find lower-quality goods gain market share in recessions. Jaimovich, Rebelo, and Wong (2019) link consumption smoothing along the quality margin to exacerbated declines in labor demand during recessions. Jørgensen and Shen (2019) document that, facing hardship, rich and middle-class households adjust consumption at the quality margin, while the poor consumers are more likely to adjust along the quantity margin.

My mechanism relies on the interplay of these two ingredients. During a recession, demand shifts toward more affordable, lower-quality goods within the same sector. Producers of these lower-quality goods, therefore, gain market share and charge higher markups. Simultaneously, firms offering higher-quality goods experience a loss of market share and decrease markups to remain competitive. I henceforth refer to this unequal markup response to recessionary changes in spending as the *markup channel*. Quantifying this markup channel, I find that observed changes in the expenditure distribution during the Great Recession predict a 6.8-percentage-point increase in the markups of lower-quality goods and a 1.8-percentage-point decrease in those of higher-quality alternatives. A typical recession's impact on product market competition thus considerably increases real consumption inequality.

Using detailed data from the NielsenIQ Homescan Consumer Panel, I show that my nonhomothetic demand system aligns closely with micro-level evidence on consumption behavior. Specifically, I establish three relevant facts that hold true in both my model and the data.

First, wealthier households allocate a larger share of their spending to more expensive, higher-quality goods. That is, households in the bottom quintile of the expenditure distribution purchase goods that are, on average, 0.56 standard deviations less expensive than what is typically charged for close substitutes. By contrast, households in the top expenditure quintile purchase goods that are roughly 0.42 standard deviations more expensive than close substitutes.

Second, I find a great deal of consumption polarization. Poor households primarily consume goods priced within an interquartile range of 0.71 to 0.34 standard deviations below the average cost of substitutes. Similarly, wealthier households purchase goods within an interquartile range of 0.16 to 0.59 standard deviations above average. By contrast, consumers in the middle expenditure quintile choose a broad mixture of cheap and expensive goods. While their typical purchase is about 0.05 standard deviations below average, their interquartile range spans from 0.36 standard deviations below to 0.45 standard deviations above average.

Third, I find that household-level price elasticities for a specific variety decrease as households increase their spending shares on that variety. For example, poorer consumers, who concentrate almost all of their spending on the least expensive options, are less sensitive to price changes in these low-cost varieties than wealthier households. This is because, in markets with limited options, a modest price increase typically does not alter which product is cheapest. As a result, poor households continue to purchase the same lowest-cost option, even when prices rise. Similarly, wealthier households exhibit little price elasticity in their purchases of premium varieties, simply because they gravitate to these higher-quality varieties. By contrast, middle-class consumers exhibit greater flexibility due to their willingness to substitute between inexpensive and expensive goods, making them comparatively responsive to price changes. This heterogeneity in price elasticities across the expenditure distribution is crucial for my mechanism; as a recession hits, households reduce spending and become locked into low-cost, lower-quality options. Producers of those goods, in turn, exploit the decrease in economy-wide price elasticity for their outputs by charging higher markups.

I calibrate my model to align with evidence on relative prices, cross-sectional differences in consumption choices, and measures of sectoral concentration. First, I set the marginal costs of different-quality goods to match the relative prices of cheap and premium goods. Next, I parameterize the nonhomotheticities in my model to reproduce spending patterns on cheap versus premium varieties. That is, conditional on matched relative prices, I target the relative expensiveness of consumption baskets across the expenditure distribution. For instance, I match the fact that purchases made by households in the top expenditure quintile are, on average, 20% more expensive than those in the bottom quintile. Additionally, I match a measure of dispersion in the expensiveness of product choices within consumer strata. I derive the expenditure distribution, a crucial determinant of these patterns, directly from Panel Study of Income Dynamics (PSID) data. Similarly, I draw sectoral heterogeneity regarding the number of different-quality producers immediately from NielsenIQ. Conditional on sector compositions, I match concen-

tration measures through the calibration of demand shifters. Finally, I calibrate the elasticity of substitution to achieve a reasonable level of aggregate markup. My model accurately reproduces a set of untargeted moments, including cross-sectional differences in price elasticities across quality tiers, as well as relative markups for cheap versus expensive varieties.

Feeding observed changes in the expenditure distribution during the Great Recession into my calibrated model, I find an unequal markup response for low- versus high-quality varieties. In particular, markups for lower-quality goods increase by an average of 6.8 percentage points, while markups for higher-quality goods decrease by 1.8 percentage points. That is, the relative price of cheaper goods increases by 5.42%. Naturally, this unequal markup response is more pronounced in markets with lower levels of competition. In sectors with a sales HHI greater than 0.35, for instance, the relative price of cheaper goods increases by 6.12%. To isolate the impact of the markup channel, my quantification exercise keeps the marginal cost of producers counterfactually fixed. In the data, where price changes are also driven by the recession's impact on marginal costs, I find the relative price of cheap versus premium goods remains countercyclical.

The markup channel is quantitatively significant in understanding the impact of a recession on consumers across the expenditure distribution. In the PSID, the Great Recession manifested as a drop in overall spending, along with a slight narrowing of the expenditure distribution. Households in the bottom expenditure quartile reduced nominal spending by 13.5% versus 16.3% in the top quartile. This reduction of inequality in nominal spending might misleadingly suggest wealthy households were more severely affected by the recession. However, accounting for the markup channel and deflating expenditures accordingly reveals consumption inequality actually widens. Real expenditures for the poor fell by 18.1% and by 15.1% among richer households. That is, poor households are disproportionately hit by recessions.

Finally, I examine the consequences of the markup channel for redistributive stabilization policy. As the markup channel increases real consumption inequality,

policymakers concerned about inequality might consider redistributive interventions. By embedding the markup channel into a full-fledged Bewley-Aiyagari-Hugget model, I analyze the effects of such redistributive measures in general equilibrium. A redistributive automatic stabilizer, which increases labor taxes by 2 percentage points for every 1% drop in aggregate TFP, significantly raises the relative price of lower-quality goods during a recession, more than doubling the increase relative to a scenario without redistribution. Intuitively, when resources are redistributed to the poor, funds that would have been spent on high-quality goods are redirected to spending on lower-quality goods. Consequently, producers of lower-quality goods capture even more market share and raise their markups further. While the markup channel itself increases inequality, redistribution amplifies its effects.

Related Literature. Cementing the premise of my paper, a large body of literature studies how households adjust their consumption during recessions. Bils and Klenow (1998) show expenditures on luxuries are substantially more cyclical. Other seminal contributions along those lines include Browning and Crossley (2000), Aït-Sahalia, Parker, and Yogo (2004) and, more recently, Jaimovich, Rebelo, and Wong (2019) and Andreolli, Rickard, and Surico (2024). Burstein, Eichenbaum, and Rebelo (2005) show lower-quality goods gain market share in recessions. Jørgensen and Shen (2019) document that rich and middle-class households smooth consumption along the quality margin, while poor consumers are more likely to adjust at the quantity margin.

Complementing my paper, a burgeoning literature studies how markups, prices, and inflation rates differ across the income distribution. Sangani (2024) documents rich households pay significantly higher retail markups for the same barcode. Here, differences in markups are due to differences in search behavior rather than product choice. Building on Kaplan and Menzio (2016), Nord (2024) connects retail price dispersion with search effort across the expenditure distribution. The response of low-quality markups to drops in spending is diametrically opposed to my predictions. As a recession hits, households throughout the economy spend

less and, in turn, perceive search as less burdensome such that retailers lose market power and charge lower markups across all quality tiers. This is complementary to my argument in that I abstract away search behavior and assume perfectly competitive retailers, whereas Nord (2024) abstracts from producer competition. Kaplan and Schulhofer-Wohl (2017) and Jaravel (2019) document poor consumers experience higher inflation rates. My findings suggest these patterns are mostly driven by the sustained price impact of contractionary episodes.

Lastly, my paper contributes to the modeling of nonhomotheticities. Seminal works include Kongsamut, Rebelo, and Xie (2001), Buera and Kaboski (2009), Boppart (2014), and Comin, Lashkari, and Mestieri (2021). With the recent and independent exception of Mongey and Waugh (2024), nonhomotheticities typically do not influence product market competition. By contrast, my framework features oligopolistically competitive firms whose strategic interactions are shaped by nonhomotheticities.

2 Model

In this section, I present a static, partial equilibrium model in which markups respond to changes in the expenditure distribution. My mechanism hinges on two critical features: nonhomothetic preferences over varieties and oligopolistic product market competition.

First, I explain how strategic firm interactions pin down markups with a general nonhomothetic demand framework. Next, I specialize the demand system to a particular class of nested nonhomothetic CES preferences that align closely with the micro data. Finally, I discuss how the implied price elasticities shape markups.

2.1 Markups under nonhomothetic demand

Firms produce varieties of differing quality and compete in oligopolistic markets. Households vary in their quality choices based on their spending, creating distinct

customer compositions for producers of low- versus high-quality goods. Moreover, since households systematically differ in their price elasticities vis-à-vis different-quality varieties, the composition of customer base matters for market power and, therefore, markups.

Environment. There is a continuum of sectors $s \in \mathscr{S}$, each containing a finite number of quality bins $q \in \{1, \ldots, Q\}$. Within each quality bin, there are a finite number of producers $i \in \{1, \ldots, N_{qs}\}$. Each producer markets a single variety (i, q, s) and operates under constant marginal cost λ_{iqs}.

The economy is also populated by a continuum of consumers who differ in their expenditure levels y. The expenditure distribution is exogenously given and characterized by a density $g(y)$. Consumer behavior is described by general nonhomothetic Marshallian demand functions $c_{iqs}(y, \boldsymbol{p})$. Following this general discussion, I will derive a specific Marshallian demand system from a class of nonhomothetic preferences, which I later demonstrate aligns closely with the data.

Firm profits. Under Bertrand competition, firms set prices to maximize profits taking as given their competitors' prices $\boldsymbol{p}_{-iqs}$, their customers' demand functions $c_{iqs}(y, \boldsymbol{p})$, and the exogenous expenditure distribution $g(y)$. Firm profits are

$$\pi_{iqs}\big(\boldsymbol{p}, g; \lambda_{iqs}\big) = \int c_{iqs}(y, \boldsymbol{p})\big(p_{iqs} - \lambda_{iqs}\big)\, g(y)\, dy. \tag{1}$$

Note that preferences are homothetic only if $c_{iqs}(y, \boldsymbol{p})$ is linearly homogeneous in y. Consequently, with homothetic preferences, profits in (1) merely scale with aggregate expenditures, and the expenditure distribution is immaterial for the producers' profit maximization problem. By contrast, with nonhomothetic preferences, strategic firm interactions are shaped by the expenditure distribution.

Customer base. With nonhomothetic preferences, households differ in their consumption choices along the quality margin based on their spending. This variation leads producers of low- versus high-quality goods to face distinct compositions of their customer base. The consumption of (i, q, s) of a household with expenditures

y relative to the aggregate consumption of (i, q, s) is denoted by

$$\tilde{c}_{iqs}(y, \boldsymbol{p}, g) \equiv \frac{c_{iqs}(y, \boldsymbol{p})}{\int c_{iqs}(y, \boldsymbol{p})\, g(y)\, dy}. \tag{2}$$

This expression provides a measure of the relative importance of consumers with expenditures y for the customer base of producer (i, q, s). Note that, under homothetic preferences, the consumption of different varieties linearly scales with overall expenditure levels. That is,

$$\tilde{c}_{iqs}(y, g) = \frac{y}{\int y\, g(y)\, dy}$$

such that the customer base is homogenous across producers.

Price elasticities. With nonhomothetic preferences at the variety level, households' price elasticities for different-quality varieties also depend on spending. The price elasticity of variety (i, q, s) among consumers of type y is denoted by

$$\varepsilon_{iqs}(y, \boldsymbol{p}) \equiv \left| \frac{\partial \log c_{iqs}(y, \boldsymbol{p})}{\partial \log p_{iqs}} \right|. \tag{3}$$

With homothetic preferences we can write $c_{iqs}(y, \boldsymbol{p}) = c_{iqs}(1, \boldsymbol{p}) \cdot y$ which is tantamount to saying that price elasticities are independent of y. With nonhomothetic preferences, however, we cannot multiplicatively separate the dependence of Marshallian demand on expenditures and prices. Consequently, price elasticities differ across the expenditure distribution.

Equilibrium. The economy is represented by a marginal cost distribution $\{\lambda_{iqs}\}$, an exogenous expenditure distribution $g(y)$, and a system of Marshallian demand functions $\{(y, \boldsymbol{p}) \mapsto c_{iqs}(y, \boldsymbol{p})\}$. Firms compete in prices. The Bertrand equilibrium is then defined as a price vector $\boldsymbol{p}^* = (p^*_{iqs})$ such that consumption allocations are consistent with $\{c_{iqs}(y, \boldsymbol{p}^*)\}$ and firms' pricing decisions constitute a

Nash equilibrium. That is, $\boldsymbol{p}^*$ solves

$$\int \left(\left. \frac{\partial c_{iqs}(y, \boldsymbol{p})}{\partial p_{iqs}} \right|_{\boldsymbol{p}^*} \left(p^*_{iqs} - \lambda_{iqs} \right) + c_{iqs}(y, \boldsymbol{p}^*) \right) g(y)\, dy = 0 \quad \forall \quad (i, q, s). \tag{4}$$

Appendix B outlines an algorithm to compute this equilibrium numerically.

Markups. Firms set markups in reference to some concept of demand elasticity. With nonhomothetic demand, the relevant notion of demand elasticity for producer (i, q, s) is the cross-sectional average of consumer-specific price elasticities weighted by the corresponding relative consumption shares. That is,

$$\mathscr{E}_{iqs}(\boldsymbol{p}, g) \equiv \int \varepsilon_{iqs}(y, \boldsymbol{p})\, \tilde{c}_{iqs}(y, \boldsymbol{p}, g)\, g(y)\, dy. \tag{5}$$

Intuitively, producers consider differences in price-elasticities across the entire population, $\varepsilon_{iqs}(y, \boldsymbol{p})$, but they also take into account which consumers ultimately matter for their customer base, $\tilde{c}_{iqs}(y, \boldsymbol{p}, g)$, as well as how many of those consumers are actually present, $g(y)$.

The gross markup is defined as price over marginal cost, $\mu_{iqs} \equiv p_{iqs}/\lambda_{iqs}$. The first-order conditions in (4) dictate that profit-maximizing markups are given by the following Lerner-type formula:

$$\mu_{iqs}(\boldsymbol{p}^*, g) = \frac{\mathscr{E}_{iqs}(\boldsymbol{p}^*, g)}{\mathscr{E}_{iqs}(\boldsymbol{p}^*, g) - 1}. \tag{6}$$

The trade-offs encapsulated by this formula depend on the details of price elasticities and consumption shares, and thus on the specifics of the demand system.

2.2 Demand under variety-level nonhomotheticities

There is a continuum of consumers with nonhomothetic preferences over varieties, differing in their expenditure levels. As varieties are differentiated along a quality margin, a change in the expenditure distribution shifts spending on different-quality goods.

Preferences. Consumers choose allocations $\{c_{iqs}\}$ to maximize real consumption of a composite final good c. The aggregation of varieties into overall consumption is based on a nested nonhomothetic CES structure. At the outer nest, real consumption c is a homothetic CES aggregate of sectoral consumption c_s. Specifically, I aggregate over a continuum of sectors $\mathscr{S}$ with

$$\int_{\mathscr{S}} \left(\frac{c_s}{c}\right)^{\frac{\eta-1}{\eta}} ds = 1. \tag{7}$$

Across-sector substitutability satisfies $\eta \geq 1$ such that c_s are gross-substitutes.

At the inner nest, nonhomotheticies encode a quality distinction. As a result, varieties are not only imperfect substitutes but also asymmetrically differentiated along the quality margin. Sectoral consumption aggregates c_s are implicitly defined through

$$\sum_{q=1}^{Q}\sum_{i=1}^{N_{qs}} \psi_q(c_s)^{\frac{1}{\sigma}} \left(\frac{c_{iqs}}{c_s}\right)^{\frac{\sigma-1}{\sigma}} = 1 \quad \forall \quad s \in \mathscr{S} \tag{8}$$

where

$$\psi_q(c_s) \equiv \frac{\varphi_q}{c_s^{(\sigma-1)(\xi_q-1)}}$$

is a nonhomothetic taste shifter. I assume that $\sigma > \eta$ to ensure that consumption is more substitutable within sectors than across sectors. The specific functional form in (8) is based on the nonhomothetic CES preferences from Comin, Lashkari, and Mestieri (2021). The key novelty in my framework is that these preferences apply at the within-sector level. As a result, with a finite number of firms in each sector, nonhomotheticities affect strategic firm interactions.

The parameters φ_q reflect a "consensus" on product quality, while the nonhomotheticity parameters ξ_q govern cross-sectional differences in quality appreciation. Specifically, φ_q acts as a demand shifter that is homogenous across the expenditure distribution. *Ceteris paribus*, when increasing φ_q for quality bin q, households uniformly shift spending toward varieties in this particular quality bin,

irrespective of consumption levels. A key feature of nonhomothetic preferences, however, is that the appreciation of quality depends on consumption. To capture this formally, the nonhomothetic demand shifter $\psi_q(c_s)$ depends on sectoral consumption c_s. Specifically, with gross substitutes, $\psi_q(c_s)$ is strictly monotonically increasing in c_s iff $\xi_q < 1$. As a result, rich households with high consumption levels spend relatively more on low-ξ varieties, whereas poor households gravitate towards high-ξ varieties. Note that when setting $\xi_q = 1$ for all q, equations (7) and (8) specialize to the familiar homothetic nested CES structure from Atkeson and Burstein (2008).

Demand for varieties. The preferences in (7) and (8) provide markets with a demand structure. Although Marshallian demand functions are not available in closed form, the nonhomothetic CES structure allows for a great deal of characterization in terms of sharp analytical expressions.

The consumers' optimization problem is best approached in two steps. First, I focus on the consumers' within-sector expenditure minimization. In each sector s, for a given price vector $\boldsymbol{p}_s = \big(p_{iqs} : i, q \big)$, the Hicksian demand to attain aggregate sectoral consumption c_s solves

$$\min_{\{c_{iqs}\}} \left\{ \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} p_{iqs} c_{iqs} \;\middle|\; \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} \varphi_q^{\frac{1}{\sigma}} \left(\frac{c_{iqs}}{c_s^{\xi_q}} \right)^{\frac{\sigma-1}{\sigma}} = 1 \right\}. \tag{9}$$

The solution to this consumer program is

$$c_{iqs}(c_s, \boldsymbol{p}_s) = \psi_q(c_s) \left(\frac{p_{iqs}}{p_s(c_s, \boldsymbol{p}_s)} \right)^{-\sigma} c_s \tag{10}$$

where the nonhomothetic ideal price index is given by

$$p_s(c_s, \boldsymbol{p}_s) \equiv \left(\sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} \tilde{p}_{iqs}(c_s)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \quad \text{and} \quad \tilde{p}_{iqs}(c_s) \equiv \psi_q(c_s)^{\frac{1}{1-\sigma}} \, p_{iqs}.$$

Intuitively, we think of $\tilde{p}_{iqs}(c_s)$ as a quality-adjusted price. With nonhomothetic preferences, the appreciation of quality depends on consumption levels; therefore,

the nonhomothetic ideal price index is also affected by consumption. Note that under homothetic CES preferences, with $\xi_q = \xi$ for all q, the ideal price index is homogeneous in sectoral consumption and Hicksian demand is linear in c_s.

Demand for sectoral aggregates. The quality distinction at the inner nest complicates the expenditure-minimizing choice of sectoral consumption. Consumers internalize the effect their allocations have on their nonhomothetic sectoral price indices. Taking as given the price vector $\boldsymbol{p} = \left(\boldsymbol{p}_s : s \right)$, the Hicksian demand for sectoral aggregates to attain an overall consumption of c solves

$$\inf_{\{c_s\}} \left\{ \int_{\mathscr{S}} p_s(c_s, \boldsymbol{p}_s)\, c_s\, ds \,\middle|\, \int_{\mathscr{S}} \left(\frac{c_s}{c} \right)^{\frac{\eta-1}{\eta}} ds = 1 \right\}. \tag{11}$$

Since the nonhomothetic ideal price index depends on c_s, this is akin to a homothetic expenditure-minimization problem with a non-linear pricing structure. Envisioning the continuum of sectors as a large set of cardinality S, as is common in Atkeson and Burstein (2008) settings, the corresponding first-order conditions dictate that

$$\left(\frac{c_s}{c} \right)^{\frac{\eta-1}{\eta}} = \frac{\sum_q \sum_i p_{iqs}\, c_{iqs}(c_s, \boldsymbol{p}_s)\, \xi_q}{\sum_s \sum_q \sum_i p_{iqs}\, c_{iqs}(c_s, \boldsymbol{p}_s)\, \xi_q} \quad \forall \quad s = 1, \ldots, S. \tag{12}$$

For each desired level of real consumption $c \in \mathbb{R}_+$, equations (12) are a set of S non-linear equations in S unknowns that pin down the Hicksian demand for sectoral consumption $c_s(c, \boldsymbol{p})$. Note that under homothetic CES preferences, consumers simply equate the left-hand side expression in (12) with their corresponding sectoral expenditure share.

Marshallian demand. The consumers' Marshallian demand allocates varieties $\left\{ c_{iqs} \right\}$ to maximize the utility from composite real consumption c for a given budget y. Formally, the Marshallian demand functions of a consumer of type y

solve

$$\arg\sup_{\{c_{iqs}\}} \left\{ c \;\middle|\; \int_{\mathscr{S}} \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} p_{iqs} c_{iqs}\, ds \leq y \text{ and aggregators (7) and (8)} \right\}. \quad (13)$$

By duality, Hicksian demand translates into Marshallian demand. Here, "indirect" real consumption $c(y, \boldsymbol{p})$ satisfies

$$\int_{\mathscr{S}} p_s\big(c_s(c, \boldsymbol{p}), \boldsymbol{p}_s\big)\, c_s(c, \boldsymbol{p})\, ds = y. \quad (14)$$

The corresponding nonhomothetic ideal price index is then defined as

$$p(y, \boldsymbol{p}) \equiv \frac{y}{c(y, \boldsymbol{p})}. \quad (15)$$

With homothetic preferences, $c(y, \boldsymbol{p})$ is linear in y, and the ideal price index is constant across the expenditure distribution. In a slight abuse of notation, the Marshallian demand for variety (i, q, s) is henceforth denoted by

$$c_{iqs}(y, \boldsymbol{p}) = c_{iqs}\Big(c_s\big(c(y, \boldsymbol{p}), \boldsymbol{p}\big), \boldsymbol{p}_s\Big) \quad \forall \quad (i, q, s). \quad (16)$$

With nonhomothetic preferences, the properties of these demand functions differ across the expenditure distribution. To build intuition, the subsequent paragraphs discuss these properties.

Expenditure elasticities. Households' expenditure elasticities determine how households differ in their quality choices and, therefore, how the expenditure distribution impacts demand patterns along the quality margin. Specifically, we can examine how sectoral expenditure shares for different-ξ varieties move with sectoral consumption levels. From equation (10), Hicksian expenditure shares are given as

$$x_{iqs}(c_s, \boldsymbol{p}_s) \equiv \frac{p_{iqs}\, c_{iqs}(c_s, \boldsymbol{p}_s)}{p_s(c_s, \boldsymbol{p}_s)\, c_s} = \psi_q(c_s) \left(\frac{p_{iqs}}{p_s(c_s, \boldsymbol{p}_s)} \right)^{1-\sigma} \quad (17)$$

and depend on c_s through both the nonhomothetic demand shifter $\psi_q(c_s)$ and the nonhomothetic ideal price index $p_s(c_s, \boldsymbol{p}_s)$. We can naturally think of a variety as being of higher quality iff its expenditure share increases in sectoral real consumption. The elasticity of x_{iqs} with respect to c_s is given as

$$\frac{\partial \log x_{iqs}(c_s, \boldsymbol{p}_s)}{\partial \log c_s} = \Big(\sigma - 1\Big)\Big(\bar{\xi}_s(c_s, \boldsymbol{p}_s) - \xi_q\Big) \tag{18}$$

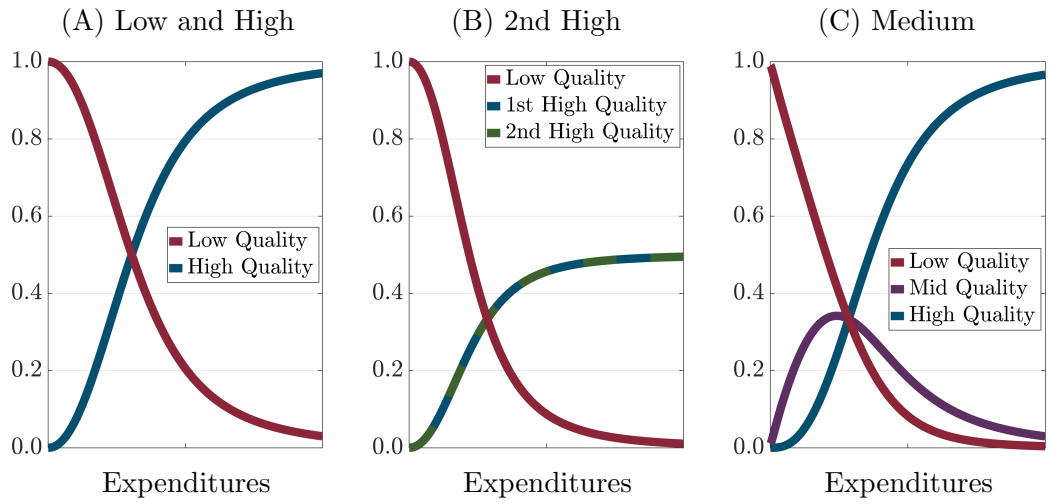
where

$$\bar{\xi}_s(c_s, \boldsymbol{p}) \equiv \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} x_{iqs}(c_s, \boldsymbol{p}_s)\, \xi_q.$$

Note that $\bar{\xi}_s$ is the average nonhomotheticity parameter for a consumer with sectoral consumption c_s. With gross-substitutes, a particular household's expenditure share on varieties in quality bin q increases in c_s iff ξ_q is below this household's average nonhomotheticity parameter. It follows that the lowest-ξ variety is unambiguously perceived as being of high quality and vice versa. Since $\lim_{c_s \to 0} \bar{\xi}_s(c_s, \boldsymbol{p}) = \max\{\xi_q\}$, poor consumers tend to perceive mid-ξ varieties as being of high quality, while richer household, for whom $\bar{\xi}_s \to \min\{\xi_q\}$, view the exact same varieties as inferior. Generally, for $Q > 2$ and $q \notin \partial\mathscr{Q}$, quality is not an intrinsic feature of a variety but rather a matter of perception, which is contingent on consumption and, therefore, ultimately expenditure levels.

Figure 1 illustrates expenditure shares as a function of expenditures. Panel A shows these Engel curves for a sector with $Q = 2$ and $N_{qs} = 1$. As spending increases, consumers allocate a greater portion of their budget to the high-quality (low-ξ) variety. Panel B introduces a second high-quality option. Here, poor consumers' spending remains concentrated on the low-quality (high-ξ) option, while rich households divide their spending between the two high-quality varieties. Panel C, introduces a mid-ξ variety. As poor consumers spend more, they allocate a larger share of their spending to this mid-ξ variety, which they perceive as high quality. Conversely, more affluent consumers decrease their relative spending on what they now view as an inferior product.

Figure 1: Expenditure Shares as a Function of Expenditures



Panel A depicts within-sector expenditure shares for a sector with a single high-ξ (low-quality) and a single low-ξ (high-quality) variety. Panel B adds a second low-ξ (high-quality) variety, whereas panel C adds a mid-ξ (medium-quality) variety. In panel B, with multiple low-ξ varieties, rich households divide there spending among those high-quality options, while poor households continue to purchase almost exclusively low-quality goods. In panel C, medium-quality goods are predominantly consumed by middle-class households.

Price elasticities. Households' price elasticities play a crucial role in determining how demand responds to price changes and are thus directly linked to firms' market power. With nonhomothetic preferences, these elasticities differ across households based on their expenditure levels.

Under Bertrand competition, the price elasticity of variety (i, q, s) among consumers with expenditures y is given as

$$\varepsilon_{iqs}(y, \boldsymbol{p}) = \Big(1 - x_{iqs}(y, \boldsymbol{p})\Big)\, \sigma + x_{iqs}(y, \boldsymbol{p})\ \eta\ \zeta_{qs}(y, \boldsymbol{p}) \tag{19}$$

where

$$\zeta_{qs}(y, \boldsymbol{p}) \equiv \frac{\Big(\sigma\, \bar{\xi}_s(y, \boldsymbol{p}) + (1 - \sigma)\, \xi_q\Big)^2}{\sigma\, \eta\ \bar{\xi}_s(y, \boldsymbol{p})^2 + (1 - \sigma)\, \eta\ \bar{\xi}_s^2(y, \boldsymbol{p}) + (1 - \eta)\, \bar{\xi}_s(y, \boldsymbol{p})}$$

and

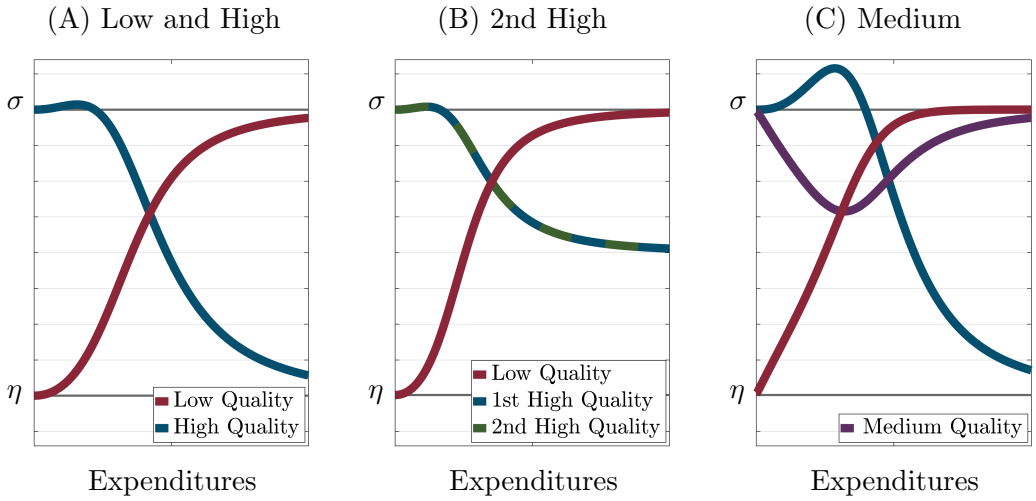
$$\bar{\xi}_s^2(c_s, \boldsymbol{p}) \equiv \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} x_{iqs}(y, \boldsymbol{p})\, \xi_q^2.$$

The price elasticity in equation (19) represents a convex combination of the within-sector elasticity of substitution (σ) and a modified form of the across-sector elasticity of substitution ($\eta \times \zeta_{qs}$). The additional term $\zeta_{qs}(y, \boldsymbol{p})$ captures that prices do not only influence sectoral price indices but, through their impact on c_s, also affect how households perceive quality. Thus, $\zeta_{qs}(y, \boldsymbol{p})$ is a reflection of the "non-linearity" in $p_s(c_s, \boldsymbol{p}_s)$. Note that ζ_{qs} depends on $(y, \boldsymbol{p})$ only through expenditure shares. Consequently, cross-sectional heterogeneity in price elasticities is fully explained by x_{iqs}.

For $\eta \approx 1$, which is consistent with the data, we have $\lim_{y\to 0} \zeta_{qs}\, x_{iqs} \approx \lim_{y\to 0} x_{iqs}$ and $\lim_{y\to\infty} \zeta_{qs}\, x_{iqs} \approx \lim_{y\to\infty} x_{iqs}$. Therefore, and since $\sigma > \eta$, the key insight from equation (19) is that the larger the expenditure share a particular household allocates to a specific variety (i, q, s), the less price-elastic they are regarding that variety. For instance, in sectors with a single inexpensive low-quality variety, consumers with $y \to 0$ allocate almost 100% of their spending to this option. Since they do not view pricier high-quality varieties within the same sector as feasible substitutes, their price elasticity approaches η. From the perspective of these consumers, there is effectively no within-sector competition, but only across-sector competition with other low-quality varieties in different sectors. This highlights the importance of competition within quality bins: when poor households can distribute their spending over multiple low-quality options, their weight on η decreases.

Similarly, rich consumers gravitate toward pricier, higher-quality products without significant regard for price. In markets with a single high-quality option, they do not even consider substitution for lower-quality alternatives. In sectors with multiple high-quality option, however, they recognize their substitutability and are thus more responsive when the price of one of those goods changes. More interestingly, middle-class households, which consume a mixture of low- and high-quality goods, are, in principle, willing to substitute along the quality margin and, therefore, comparatively price-elastic in either direction. This holds true even in relatively concentrated sectors.

Figure 2: Price-Elasticities as a Function of Expenditures



Panel A depicts price-elasticities for a sector with a single high-ξ (low-quality) and a single low-ξ (high-quality) variety. Panel B adds a second low-ξ (high-quality) variety, whereas panel C adds a mid-ξ (medium-quality) variety. In panel A, poor and rich households, whose consumption is concentrated, are price-inelastic regarding their preferred ξ. Middle-class households are price-elastic in either direction. Adding a second high-quality variety in panel B means that rich consumers have more options and are thus more price-elastic vis-à-vis high-quality goods. These options are not feasible for poor households and their price-elasticities remain unchanged. In panel C, the addition of a medium-quality good increases price-elasticities for middle-class households. Poor and rich households remain price-inelastic for the bulk of their consumption.

Figure 2 illustrates price elasticities as a function of expenditure. Panel A shows price elasticities for $Q = 1$ and $N_{qs} = 1$. Panel B depicts a less concentrated sector with two high-quality options. Here, poor consumers remain price-inelastic for the low-quality option, as the additional pricier high-quality variety does not qualify as an affordable substitute. By contrast, with more options in the high-quality segment, richer consumers become more price-responsive relative to panel A. Panel C introduces an additional medium-quality variety, which is primarily consumed by middle-class households. This additional option increases middle-class price elasticities across the board.

Note that the price elasticity in (19) generalizes a more familiar setting. When $\xi_q = 1$ for all q, preferences specialize to a nested homothetic CES structure. In that case, expenditure shares are constant across the expenditure distribution, and $\zeta_{qs}(y, \boldsymbol{p})$ is identically equal to one. Consequently, the price elasticity in (19)

collapses into the standard expression from Atkeson and Burstein (2008)

$$\varepsilon_{iqs}(\boldsymbol{p}) = \Big(1 - x_{iqs}(\boldsymbol{p})\Big)\,\sigma + x_{iqs}(\boldsymbol{p})\,\eta.$$

Moreover, if there is no distinction between within- and across-sector substitutability ($\sigma = \eta$) we eliminate the effects of granularity and the price elasticity even further specializes to the expression obtained under monopolistic competition, $\varepsilon_{iqs} = \sigma$.

Markups. The demand elasticity $\mathscr{E}_{iqs}(\boldsymbol{p}, g)$ in equation (5) is given as the cross-sectionally averaged consumer-specific price elasticity ε_{iqs} weighted by relative consumption shares $\tilde{c}_{iqs}$. From Figure 1, we have seen that consumers in the tails of $g(y)$ concentrate their spending on either low-quality or high-quality goods, which makes them relatively price inelastic for most of their purchases. By contrast, middle-class households consume a mixture of low-, medium-, and high-ξ varieties. With this greater willingness to substitute along the quality margin, they have effectively more options and are therefore comparatively price elastic.

When setting markups, firms trade off the loss of business from relatively price-elastic middle-class customers against the rents they could extract from their less elastic customer segments. Consider, say, a producer selling an expensive, high-quality variety. Their customer base consists of price-elastic middle-class households and extremely inelastic consumers in the right tail of $g(y)$. When considering a price increase, this producer weighs the loss of business from middle-class consumers against the higher margin earned from their price-insensitive affluent customer segment. This trade-off accounts for the mass of each consumer type within the population. Strategic price-setting, therefore, depends on the expenditure distribution.

3 Micro evidence for nonhomothetic demand

In this section, I demonstrate that the key predictions of my demand system align well with empirical observations on consumption choices across the expenditure distribution. These predictions are crucial for connecting recessionary drops in spending to shifts in demand patterns that ultimately lead to an unequal markup response. Specifically, I find that rich households spend relatively more on premium goods and that there is substantial consumption polarization across the expenditure distribution. Moreover, the data reveal varying price elasticities for different-quality varieties across consumer strata.

Data sources. My primary dataset for this exercise is the *NielsenIQ Homescan Panel*. This dataset, made available by the Chicago Booth Kilts Center for Marketing Research, tracks the shopping behavior of approximately 50,000 U.S. households from 2004 to 2022. It consists of unbalanced longitudinal data on barcode-level quantities and prices of fast-moving consumer goods purchased from a wide range of retail outlets across the U.S. The data are projectable to the entire U.S. Panelists use in-home scanners to record their purchases intended for personal use. The dataset includes a wide array of self-reported demographics. Overall, the universe of NielsenIQ barcodes accounts for about 30-40% of spending on goods and roughly 15% of total expenditures.

I supplement my empirical analysis with *NielsenIQ Retail-Scanner* data, which provides weekly pricing, volume, and store information generated by point-of-sale systems from over 90 participating retail chains across all U.S. markets. This dataset covers scanner data from 35,000 to 50,000 participating grocery and drug stores, accounting for more than half of the total sales volume in the U.S.

Furthermore, I use wholesale data from *PriceTrak PromoData* to correlate retail markups with barcode expensiveness. PromoData is a weekly monitoring service that tracks wholesale prices for a subset of NielsenIQ barcodes. The data are sourced from 12 grocery wholesaler organizations that sell products to retailers across the U.S. and covers the period from 2006 to 2012.

A premium margin in the data. Households with different expenditure levels vary in their propensities to consume cheap versus expensive goods. To define a meaningful measure of expensiveness in the data, I first demarcate a set of close substitutes for each barcode i. Here, I rely on NielsenIQ's classification of barcodes into narrowly defined product modules m. These product modules (e.g. fresh apples, mozzarella cheese, or instant coffee) consist of highly substitutable barcodes. To ensure within-module comparability of price data, I convert barcode-level quantities and prices to module-specific base units. Whenever a natural conversion is not feasible (e.g. count vs. ounces), modules are segmented accordingly.

To define a barcode-level measure of expensiveness, I proceed as follows: For each region r and year t, I compute the average price per base unit of each barcode i based on regular price data. Since I average over households h and different stores within region r, going forward, prices are barcode-level objects. That is,

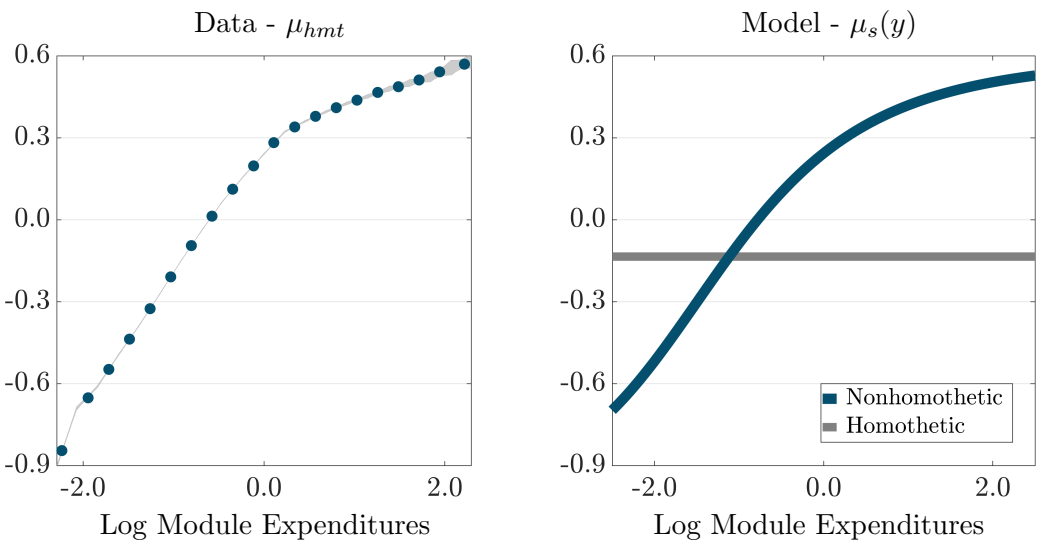
$$\text{price}_{irt} = \frac{\sum_{h \in r} \text{expenditure}_{iht}}{\sum_{h \in r} \text{quantity}_{iht}}. \tag{20}$$

In within-module comparisons, this price_{irt} effectively encodes the expensiveness of barcode i in the region-time cell (r, t). To make the notion of expensiveness comparable across modules, I residualize price_{irt} on module, region, and time fixed effects, as well as interactions thereof. I then normalize it by a module-clustered measure of residual dispersion. That is, I define

$$\text{premium}_{irt} = \frac{\text{price}_{irt} - \alpha_{\text{module}} - \alpha_{\text{region}} - \alpha_{\text{module}\times\text{region}} - \alpha_{\text{year}}}{\sigma_{\text{module}}}. \tag{21}$$

I refer to this object as a barcode-level premium score. In particular, premium_{irt} measures how many standard deviations barcode i is priced above what is typical for the corresponding module in region r at time t. Besides facilitating comparability across modules, this premium score rids the data of regional heterogeneity in product availability, pricing patterns, and overall time trends.

Figure 3: Taste for Premium Goods Increases with Expenditures



The lefthandside panel depicts a binscatter of the logarithm of region/time normalized household module-expenditures on the x-axis against household premium-indices μ_{hmt} on the y-axis. The construction of confidence band for this binscatter follows Cattaneo *et al.* (2023). The righthandside panel plots the model-implied relationship between expenditures and household premium indices under homothetic as well as nonhomothetic preferences.

Fact 1: Rich households spend more on premium goods. Households with higher expenditures spend relatively more on pricier items, even when lower-cost substitutes are available. To establish this fact in the data, I correlate consumption patterns along the premium margin with household expenditures. Specifically, I define the household premium index μ_{hmt} as a household- and module-specific quantity-weighted average of the barcode-level premium scores defined in (21). That is,

$$\mu_{hmt} = \sum_{i \in m} \frac{\text{quantity}_{iht}}{\sum_{i \in m} \text{quantity}_{iht}} \text{premium}_{ir(h)t}. \tag{22}$$

Crucially, μ_{hmt} depends on consumer behavior only through quantity choices; the price component of μ_{hmt} does not reflect search behavior or store characteristics. This premium index indicates the extent to which household h consumes cheap versus expensive varieties in module m. To illustrate the correlation between premium indices and expenditure levels, the left-hand side panel of Figure 3 depicts

a binscatter of μ_{hmt} against total expenditures in module m. We see that households that spend particularly little tend to purchase varieties that are, on average, priced about 0.9 standard deviations below what is typical in their corresponding module. As households spend more, they do not merely purchase higher quantities of the exact same varieties, but gravitate toward more expensive options.

To compare this observation in the data with my model predictions, the model counterpart of μ_{hmt} is given by

$$\mu_s(y) = \sum_{q=1}^{Q} \sum_{i=1}^{N} \frac{c_{iqs}(y, \boldsymbol{p})}{\sum_{q=1}^{Q} \sum_{i=1}^{N} c_{iqs}(y, \boldsymbol{p})} \, p_{iqs}. \tag{23}$$

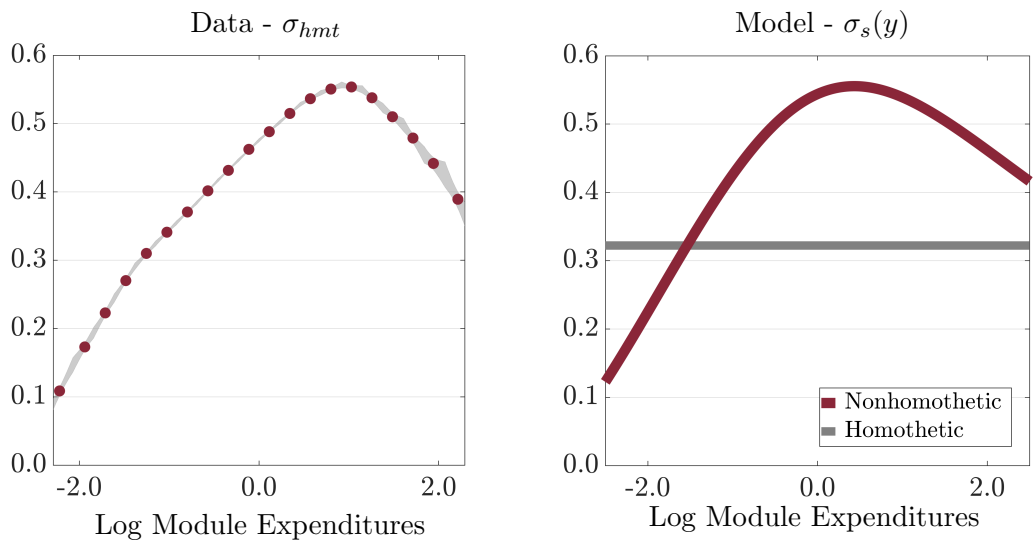
The right-hand side panel of Figure 3 shows this model-implied household premium index under homothetic versus nonhomothetic preferences. With homothetic preferences, the consumption of different varieties simply scales with expenditure levels, and the quantity weights in (23) are constant across the expenditure distribution. By contrast, with nonhomothetic preferences, the composition of consumption baskets depends on overall expenditure levels. As a result, nonhomothetic preferences are well-suited to capture the observed tendency of affluent households to spend relatively more on pricier goods.

Fact 2: Consumption polarization. I find significant polarization in consumption patterns across the expenditure distribution. Poor households almost exclusively opt for inexpensive goods, whereas wealthier households predominantly consume premium goods. By contrast, middle-class households purchase a broad mixture of varieties along the premium margin. To establish this fact in the data, I compute a measure of household premium dispersion σ_{hmt}. Specifically, this is the second moment corresponding to (22) such that

$$\sigma^2_{hmt} = \sum_{i \in m} \frac{\text{quantity}_{iht}}{\sum_{i \in m} \text{quantity}_{iht}} \Big(\text{premium}_{ir(h)t} - \mu_{hmt} \Big)^2. \tag{24}$$

Therefore, σ_{hmt} reflects the extent to which household h consumes a mixture of cheap and expensive varieties in module m; a low measure σ_{hmt} indicates that

Figure 4: Polarization in Consumption Patterns



The lefthandside panel depicts a binscatter of the logarithm of region/time normalized household module-expenditures on the x-axis against household premium-dispersion σ_{hmt} on the y-axis. The construction of confidence band for this binscatter follows Cattaneo *et al.* (2023). The righthandside panel plots the model-implied relationship between expenditures and household premium-dispersion under homothetic as well as nonhomothetic preferences.

household h's purchases consistently align with their premium index. As illustrated in the left-hand side panel of Figure 4, σ_{hmt} is comparatively small for households in either tail of the expenditure distribution. Poor households not only lean towards the consumption of cheaper goods, as shown Figure 3, but they almost exclusively purchase those inexpensive options. Similarly, richer households primarily opt for more expensive varieties. In the middle of the distribution, the middling premium index in Figure 3, however, does not reflect the exclusive consumption of typically priced goods but is due to significant mixing along the premium margin.

The model counterpart of σ^2_{hmt} is given as

$$\sigma^2_s(y) = \sum_{q=1}^{Q}\sum_{i=1}^{N} \frac{c_{iqs}(y,\boldsymbol{p})}{\sum_{q=1}^{Q}\sum_{i=1}^{N} c_{iqs}(y,\boldsymbol{p})} \Big(p_{iqs} - \mu_s(y)\Big)^2 \tag{25}$$

Since cross-sectional variation captured by $\sigma_s(y)$ is exclusively driven by differ-

ences in the composition of consumption baskets, a model with homothetic preferences cannot possibly replicate the observed pattern. The right-hand side panel of Figure 4, however, demonstrates that a nonhomothetic preference structure accurately reproduces the inverted u-shape observed in the data.

Fact 3: Households' price-elasticities decline in spending shares. I find that the well-known empirical regularity that price elasticities decline with spending shares holds true at the household level. A particular household is least price-elastic with respect to whichever variety they (individually) consume the most. Intuitively, poor consumers, who routinely buy the least expensive options, show minimal substitution responses to minor price changes for these lower-priced goods. Similarly, wealthy consumers, with a strong appetite for premium goods, exhibit very little consumption response to price fluctuations for these pricier varieties.

To establish this fact in the data, I stratify the population based on their consumption choices along the premium margin.[2] Specifically, my categorization is based on household premium indices as defined in (22). Premium consumers are those whose premium index falls within the upper tertile of the cross-sectional distribution $\{h \mapsto \mu_{hmt}\}$. Basic consumers are those with a premium index in the lower tertile. In particular, for each module m and time t, I define

$$\mathscr{H}\big|^{mt}_{\text{premium}} = \left\{ h \in \mathscr{H} \;\middle|\; \mu_{hmt} \text{ in upper tertile of } h \mapsto \mu_{hmt} \right\}$$

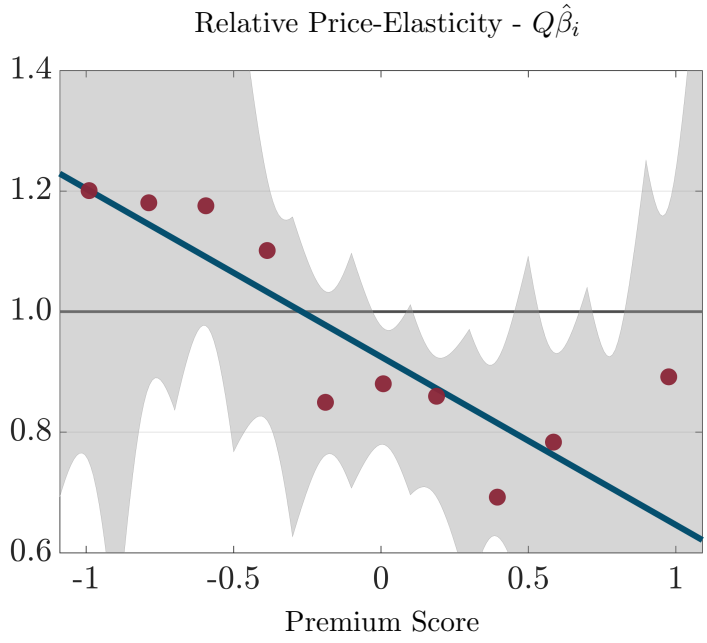
and

$$\mathscr{H}\big|^{mt}_{\text{basic}} = \left\{ h \in \mathscr{H} \;\middle|\; \mu_{hmt} \text{ in lower tertile of } h \mapsto \mu_{hmt} \right\}.$$

Then, to recover barcode-specific price elasticities for, say, premium consumers, I estimate a log-linearized version of the demand system implied by equation (16)

[2] In Appendix A, I estimate barcode-level price elasticities as a function of within-module expenditure shares. In line with Figure 5, this estimation also reveals that poor consumers exhibit lower price elasticities than rich households for lower-cost goods.

Figure 5: Differential Price Elasticities Across the Expenditure Distribution



This graph depicts a binscatter of the barcode-level premium scores on the x-axis against the corresponding relative price elasticities $Q\hat{\beta}_i$ for rich vs poor households. Confidence bands are constructed following Cattaneo, Crump, Farrell, and Feng (2023). For details see Section 3.

where an observation (i, h, t) is included iff $h \in \mathscr{H}|^{m(i)t}_{\text{premium}}$. That is,

$$\begin{aligned} \text{log quantity}_{iht} = \; & \alpha^{\text{prm}}_{ih} + \alpha^{\text{prm}}_{ir} + \alpha^{\text{prm}}_{it} + \beta^{\text{prm}}_{i} \text{log price}_{iht} \\ & + \sum\nolimits_{j \in \mathscr{K}_{iht}} \beta^{\text{prm}}_{ij} \text{log price}_{jr(h)t} + \gamma^{\text{prm}}_{i} \text{log expenditure}_{ht} + \epsilon^{\text{prm}}_{iht}. \end{aligned} \tag{26}$$

In this regression, β^{prm}_i can be interpreted as barcode i's own price elasticity among premium consumers. The regression controls for a judiciously constructed set of household-specific competitors for each barcode i. Leveraging data on shopping trips from NielsenIQ's consumer panel and store-level pricing information from NielsenIQ's retail-scanner data, I ensure that $\mathscr{K}_{iht}$ is comprised of barcodes $j \in m(i)$ that are actually available to household h at $\text{price}_{jr(h)t}$.

In order to address potential endogeneity issues in the relationship between quantity_{iht} and price_{iht}, I construct price instruments in the spirit of Hausman

(1996). Specifically, I instrument price_{iht} with the average price for barcode i in year t, excluding observations from region $r(h)$. Price elasticities are, therefore, identified by within household/region/time price variation that is explained by global price movements. Since all coefficients in (26) are indexed at the barcode level, under mild clustering conditions on $\epsilon^{\text{prm}}_{iht}$, all regressions are run barcode by barcode.

Proceeding analogously for the set of basic consumers $\mathscr{H}|^{mt}_{\text{basic}}$, I define a barcode-level measure of relative price elasticity as

$$Q\,\hat{\beta}_i \equiv \hat{\beta}^{\text{bsc}}_i \big/ \hat{\beta}^{\text{prm}}_i. \tag{27}$$

Whenever $Q\,\hat{\beta}_i > 1$, premium consumers are less price-elastic with respect to barcode i compared to basic consumers. After controlling for module fixed effects, the binscatter in Figure 5 depicts these relative price elasticities against barcode-level premium scores on the x-axis. The graph reveals that basic consumers are less price-elastic than premium consumers when it comes to inexpensive varieties. Vice versa, premium consumers are less price-elastic than basic consumers when it comes to premium items. The regularity here is that as households concentrate their spending on goods within a particular price range, they effectively encounter fewer options and are, consequently, less price-elastic; price elasticities decline in spending shares.

4 Quantification

In this section I outline my calibration strategy. I parameterize my nonhomothetic demand structure based on data moments from the NielsenIQ Homescan Consumer Panel. Additionally, I use expenditure data from the PSID to capture the distribution of household spending beyond fast-moving consumer goods.

Quantitative model. For tractability, I make a binary quality distinction with $q \in \{\text{low, high}\}$. Firms' marginal costs λ_q are quality-dependent. In my baseline

Table 1: Calibrated Parameters

Parameter	Value	Significance	Parameter	Value	Significance
Technology			Substitution		
λ_{low}	0.80	Marginal cost	σ	18	Within sector
λ_{high}	1.13	Marginal cost	η	1.02	Across sector
Quality					
$\xi_{\text{high}}/\xi_{\text{low}}$	0.74	Nonhomotheticity	φ_{low}	0.86	Demand shifter
ν	20,896	Expenditure scale	φ_{high}	1.33	Demand shifter

This table reports the internally calibrated parameters. I calibrate marginal cost for different-quality goods to match relative prices. The parameterize the nonhomothetic taste shifter in (8) to match salient reflections of facts 1 and 2 in Section 3. Shutting down nonhomotheticities, the within- and across-sector elasticities of substitution are chosen targeting moments of the model-implied markup distribution from Becker *et al.* (2024).

model there are no differences in marginal cost within quality tiers. A sector is, therefore, fully characterized by the number of firms operating in each quality bin. I partition the continuum of sectors $\mathscr{S}$ into a finite number of uncountably infinite sets of identical sectors. Specifically, I consider 25 distinct sector compositions with $\left(N_{\text{low}}, N_{\text{high}}\right) \in \left\{1, \ldots, 5\right\}^2$. The measure of each one of these compositions is given by the corresponding fraction in the data. Similarly, I read the expenditure distribution directly off of the PSID. Robustness checks with a more granular segmentation of the quality spectrum as well as productivity differences within quality bins are relegated to Appendix C.

Calibration. I calibrate the firms' marginal costs λ_q to align my model with data on relative prices along the premium margin. To recover a suitable target, I compute the sales-weighted average of below versus above median prices in each product module. I then average the resulting ratio across modules.

I calibrate the nonhomotheticity parameters $\left\{\xi_q\right\}$ to match the consumption of cheap versus expensive goods across the expenditure distribution. Since model-implied Engel curves are functions of nominal spending y, the parameterization of

Table 2: Moments Used in Calibration

Target	Source	Data	Model
Relative price (high/low)	NielsenIQ	1.25	1.24
Premium index (mid/poor)	NielsenIQ	1.06	1.07
Premium index (rich/poor)	NielsenIQ	1.21	1.20
Polarization (mid/poor)	NielsenIQ	5.04	4.48
Polarization (rich/poor)	NielsenIQ	3.18	2.41
Local sales HHI	NielsenIQ & GS1	0.23	0.23
Aggregate markup*	BEMX '24	1.31	1.32
Markup dispersion*	BEMX '24	0.23	0.19

This table reports the model fit achieved through internal calibration of the parameters in Table 1. The asterisk $*$ indicates that moments on the markup distribution from Becker, Edmond, Midrigan, and Xu (BEMX) shutting down nonhomotheticities while still matching the data moment on average local sales concentration.

$\{\xi_q\}$ is inextricably linked with scale of the expenditure distribution. Therefore, let ν denote a scalar such that $y_{\text{data}} = \nu\, y$. Conditionally on having matched relative prices, I calibrate the parameters $(\{\xi_q\}, \nu)$ targeting household premium indices and household premium dispersion across the expenditure distribution.[3]

With exogenously given sector compositions, the demand shifters $\{\varphi_q\}$ shift consumption across quality tiers and therefore speak to measures of sectoral sales concentration.[4] Since most competition is inherently local, the relevant measure of concentration is that of local sales concentration.[5] I therefore calibrate my model to match an average local sales HHI of 0.23 computed from a merge of NielsenIQ and GS1 data. This calibration target is consistent with measures of local sales concentration documented by Benkard, Yurukoglu, and Zhang (2023).

[3] As $\eta \to 1$, the scale of the nonhomotheticity parameters $\{\xi_q\}$ is increasingly less identified. Since, in my calibration $\eta = 1.025$, I normalize $\xi_{\text{low}} = 1$.

[4] In contrast to homothetic preferences, with nonhomotheticities, the distinction between φ_q and λ_q is meaningful even vis-à-vis revenue data.

[5] See e.g. Rossi-Hansberg and Hsieh (2023) or Franco (2024).

Table 3: Untargeted Moments

Moments	Data	Model
Relative price elasticity (low quality)	0.87	0.82
Relative price elasticity (high quality)	1.16	1.23
Relative retail markups	0.99	0.96

This Table reports untargeted moments. Relative price elasticities are the quotient of barcode-specific price elasticities for poor vs rich consumers. Relative retail markups are computed from PriceTrak PromoData 2006-2012. For details see Section 3.

Lastly, I temporarily shut down nonhomotheticities in my model, and pin down the parameters governing within- and across-sector substitutability to match moments of the model-implied markup distribution from Becker, Edmond, Midrigan, and Xu (2024).[6] Specifically, I choose $\{\eta, \sigma\}$ to match an aggregate markup of 1.31 in a homothetic version of my environment and for a fixed measure of local sales concentration. When reintroducing nonhomotheticities, the aggregate markup increases from 1.31 to 1.41. Intuitively, for the same level of concentration, a quality distinction means that competition is diluted. Table 2 summarizes the model fit.

Model Validation. In Table 3, I also report some key moments that were not targeted in my calibration exercise. An untargeted moment that is crucial for the response of markups to changes in spending patterns, is that of relative price elasticities for different-quality varieties. Computing the ratio of model-implied price elasticities for poor versus rich consumers in each quality bin, I find that my model is perfectly in line with the evidence presented in Section 3.

Since production markups are not readily measured in the data, I compute retail markups from PriceTrak PromoData and correlate them with barcode-level premium scores. I find no particular correlation between markups and premium

[6] As pointed out by Bond *et al.* (2022) the measurement of markups based on revenue data is an inherently bleak endeavor. I, therefore, target model-implied markups.

scores. The model-implied relative markup on high- vs low-quality goods aligns well with the data (on retail markups) and is reported in Table 3.

5 Quantification of the Markup Channel

In this section I use my model to show that, during the Great Recession, markups on lower-quality varieties increased, whereas those on higher-quality varieties fell. Recession-induced shifts in demand therefore made the Great Recession even more burdensome for poor consumers. Providing direct evidence on this mechanism, I show that relative price movements in the data are consistent with my model predictions.

5.1 An unequal markup response

To quantify the response of different-quality markups to demand shifts during the Great Recession, I feed observed changes in the expenditure distribution into the pre-crisis calibration of my model.

Changes in the expenditure distribution. The Great Recession led to a drop in overall spending alongside a slight narrowing of expenditure inequality. This observation emerges from a comparison of symmetrically PCE-deflated expenditures in PSID data for 2006 versus 2012. Given the biennial nature of the PSID, my choice of period accounts for the full impact of the crisis on consumer spending habits, acknowledging scarring effects and longer-lasting changes in expenditure patterns. While expenditures declined at the outset of the recession, the most pronounced shifts are observed when comparing pre-crisis (2006) and post-crisis (2012) periods. Expenditure levels returned to, and eventually exceeded, pre-crisis levels in subsequent years. Figure 13 in Appendix A depicts a histogram comparing the expenditure distribution in 2006 with that of 2012. There was a notable influx of expenditure mass into the lower end of the support and a slight

Table 4: Price Impact of Markup Channel

		Nonhomothetic		Homothetic	
		$\Delta\mu$	Δp	$\Delta\mu$	Δp
Overall	Low Quality	6.79 pp	4.19 %	0 pp	0 %
	High Quality	−1.82 pp	−1.21 %	0 pp	0 %
Low Competition	Low Quality	8.43 pp	3.97 %	0 pp	0 %
($\text{HHI}_{2006} \approx 0.35$)	High Quality	−2.88 pp	−1.59 %	0 pp	0 %
High Competition	Low Quality	2.57 pp	1.97 %	0 pp	0 %
($\text{HHI}_{2006} \approx 0.10$)	High Quality	−1.22 pp	−0.95 %	0 pp	0 %

This table reports the model-implied average markup and price response of low- and high-quality goods during the Great Recession. I report markup changes in terms of percentage points. Price changes are presented as percentages and can be read as percentage changes in gross markups. Results are stratified by competition levels according to pre-crisis sales HHIs. The latter columns show results for a homothetic baseline.

decrease in inequality. Specifically, average spending declined by 15.9% while the 75/25 percentile ratio of the distribution decreased from 3.02 to 2.89.

Model-implied markup response. Table 4 presents model-implied markup and price changes during the Great Recession. On average, the Great Recession caused a 6.79-percentage-point increase in model-implied markups for lower-quality varieties. By contrast, markups for high-quality goods declined by an average of 1.82 percentage points. I refer to the economic forces underlying this unequal markup response as the *markup channel*. To isolate the markup channel, in this quantitative exercise, the Great Recession manifests exclusively as a change in the expenditure distribution. Firms' marginal costs are counterfactually fixed and conform to the pre-crisis calibration of my environment.

The results in Table 4 are stratified by within-sector competition levels. Naturally, there is a more pronounced markup response in markets with lower levels of

Table 5: Impact Across the Expenditure Distribution

Quartile	Q1	Q2	Q3	Q4
Δ Nominal spending (data)	−13.5%	−12.9%	−13.7%	−16.3%
Δ Price (model)	4.5%	3.4%	0.1%	−1.2%
Δ Real spending (model)	−18.05%	−16.28%	−13.81%	−15.09%

This table quantifies the unequal impact of the Great Recession across the expenditure distribution. Changes in "nominal" spending are directly computed from (symmetrically PCE-deflated) expenditures in the PSID. Prices changes are model-implied and take into account the composition of consumption baskets in different quartiles of the expenditure distribution. Changes in real spending are computed as the difference of rows one and two.

competition, where firms capitalize on the limited options available to consumers. Conversely, in markets with high competition, the availability of multiple options curtails market power. Intuitively, the ability to substitute within quality bins means that households remain more price-elastic, even as they consume a less variegated mixture across quality tiers.

As the Great Recession caused households to spend less, there was a substantial shift of consumption toward relatively inexpensive, low-quality goods. Producers of these lower-quality varieties responded to this influx of customers by charging higher markups.[7] Specifically, I find markups for low-quality varieties increased by an average of 6.79 percentage points. With fixed marginal costs, this corresponds to a 4.19% price increase. Since even wealthier households opted for less expensive alternatives, the corresponding shift in demand patterns prompted producers of higher-quality goods to reduce markups. During the Great Recession, the markup channel led to an average 1.21% price decrease for high-quality goods.

Impact on consumers across the expenditure distribution. In PSID data, the expenditure distribution narrowed over the course of the Great Recession.

[7] As I quantify the impact of shifts in the expenditure distribution at business cycle frequencies, I deliberately abstract from entry and exit. The relative profitability of low-quality goods, therefore, does not precipitate a surge in competition in this market segment.

However, this decrease in inequality is reversed when accounting for the markup channel and deflating household expenditures accordingly. Table 5 presents results for quartiles of the expenditure distribution.

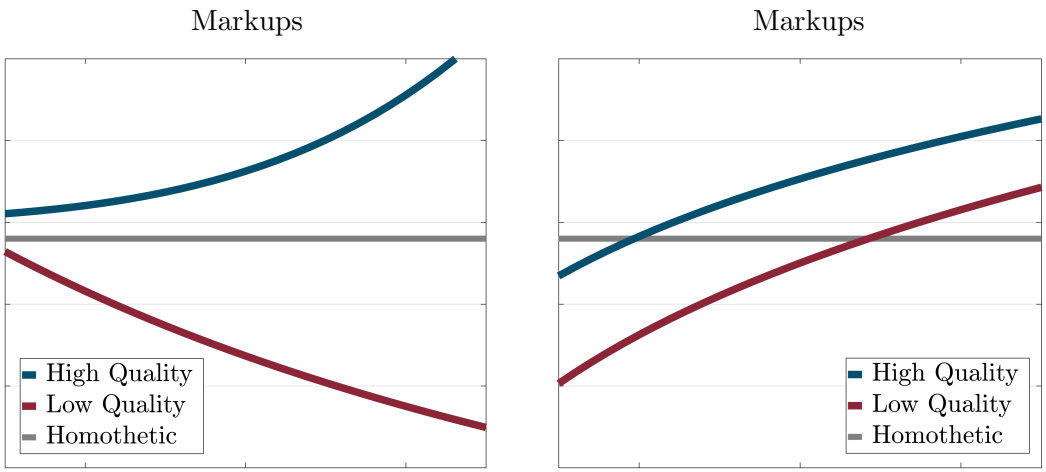
It is well known, e.g. from Heathcote *et al.* (2020), that poor households are disproportionately hit by recessions. They are more likely to lose their jobs and experience, on average, a more substantive decrease in labor earnings relative to richer households. In PSID data, the income distribution's 75/25 percentile ratio rose from 4.48 to 5.61 during the Great Recession. From the first row of Table 5, this does not seem to be true for spending at first glance. In percentage terms, the reduction in nominal expenditures among poor consumers was less pronounced than among richer households. This pattern is not implausible: affluent households can cut their spending by smoothing along the quality margin – a tool that is not available to poor households whose consumption baskets are largely comprised of low-quality options to begin with.[8] A conclusion to the effect that the Great Recession had a smaller impact on the consumption of poor households would, however, be incorrect. Since recessionary shifts in demand patterns also led to an increase in the relative price of low-quality goods, there was a force that exacerbated inequality in real terms. Accounting for the markup channel during the Great Recession, consumption among poor households actually decreased more than consumption among the rich. In real terms, consumption inequality, measured by the 75/25 percentile ratio, rose from 3.02 to 3.11.

5.2 Inspecting the mechanism

In this subsection, I provide intuition for the forces shaping this unequal markup response. Abstracting from higher-order features of the expenditure distribution, I use my model as a laboratory to examine the markup channel separately in two distinct scenarios: first, I consider a drop in overall spending; and, second, I examine a narrowing of expenditure inequality.

[8] Poor households' consumption is also closer to subsistence levels. As a result, there is not much scope to reduce spending substantialy.

Figure 6: Isolated Markup Response to Aggregate Spending & Inequality



The left-hand side panel plots model-implied markups as a function average spending levels. Entertaining a dispersion-preserving shift, the standard deviation of the expenditure distribution is kept at a fixed level. Conversely, the right-hand side panel illustrates the markup response to a change in expenditure inequality. That is, I graph markups as a function of the standard deviation of the expenditure distribution while maintaining a fixed level of aggregate spending. Under homothetic preferences markups are unresponsive to changes in the expenditure distribution.

Aggregate spending. As spending falls throughout the economy, markups along the quality margin respond asymmetrically. Specifically, markups for low-quality goods increase, while those for high-quality goods fall. The experiment I conduct here is what I call a *dispersion-preserving shift*: I vary aggregate spending while maintaining a fixed level of inequality. For clarity, I carry out this experiment in a stylized environment with a binary quality distinction and a single variety marketed in each quality bin. A drop in overall spending means that consumers across the expenditure distribution shift to more affordable, lower-quality options. Consequently, as economy-wide spending shares on low-quality goods increase, producers of these varieties gain more market power and charge higher markups. By contrast, as richer households cut their spending, they substitute toward low-quality varieties. The corresponding loss of business prompts producers of high-quality varieties to reduce markups in order to remain competitive. The lefthandside panel of Figure 6 illustrates.

Expenditure inequality. In isolation, a narrowing expenditure distribution leads to a decrease in markups across all quality bins. The experiment here is a *mean-preserving spread* I vary the standard deviation of $g(y)$ while maintaining a constant level of aggregate spending. As inequality declines, there is a shift of consumers from either tail into the middle of the distribution. Consequently, with a larger mass of middle-class consumers, there is weaker segmentation of the consumer base. As the disciplining influence of comparatively price-elastic middle-class households strengthens, firms across different quality bins engage in fiercer competition and charge lower markups. At business cycle frequencies, narrowing inequality, therefore, mitigates the adverse consequences of a recession. The righthandside panel of Figure 6 illustrates.

A remark is in order: Aguiar and Bils (2015) document a secular trend of increasing expenditure inequality over the last decades. With a correspondingly dwindling middle class, through the lens of my model, firms became less concerned about losing business from price-elastic middle-class consumers. Producers at either end of the quality spectrum, in turn, focused on extracting rents and charged higher markups. That said, a proper examination of this potential link between increasing expenditure inequality and rising markups and falling labor shares is beyond the scope of this paper. It would certainly require a model that endogenizes market structure and accommodates evidence on the evolution of entry barriers, as discussed by Philippon and Gutiérrez (2019).

A decomposition. A typical recession unambiguously places downward pressure on high-quality markups. The equilibrium response of low-quality markups, on the other hand, is shaped by two diametrically opposed forces. Table 6 presents a decomposition of the overall markup response shown in Table 4, separating the impact of the drop in spending as well as the decrease in inequality during the Great Recession. Absent the narrowing of the expenditure distribution, markups for low-quality varieties would have increased by 7.13 percentage points.

Table 6: Decomposition of the Markup Response

	Overall response	Drop in spending	Inequality	Higher moments
Low Quality	6.79 pp	7.13 pp	−2.95 pp	2.61 pp
High Quality	−1.82 pp	−1.91 pp	−0.34 pp	0.43 pp

This Table presents results for a decomposition exercise. For the 2006 baseline I assume a log-linear expenditure distribution which is parameterized to match moments of the (ν-scaled) empirical mass function. I then reparameterize to match the overall markup response from Table 4. For the decomposition exercise, I, separately, fix average spending and standard deviation while matching the respective other moment with its 2012 value. The effect of higher moments, interactions, as well as approximation error is reported in the last column.

5.3 Evidence on the mechanism

In this subsection, I show that, in the data, both relative price movements and shifts in spending patterns during the Great Recession are consistent with my model predictions. By attributing quantitative discrepancies to other forces, such as recessionary changes in marginal costs, I assess the relative importance of the markup channel.

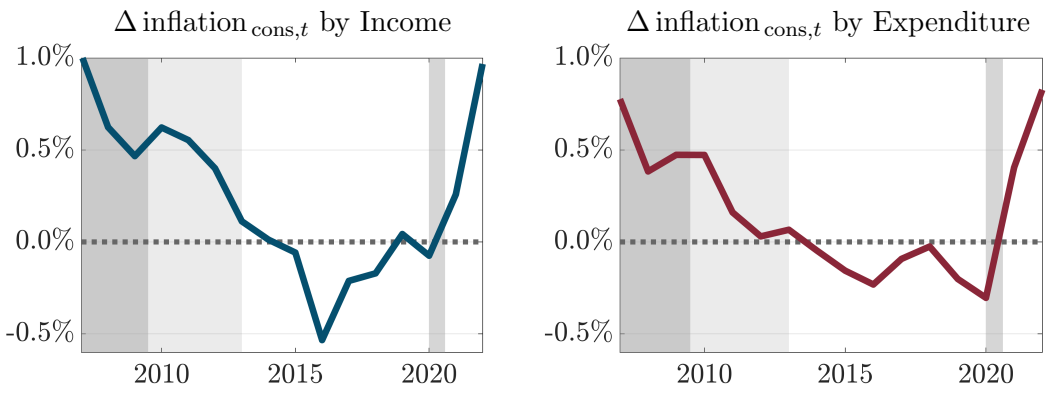
Differences in inflation rates. Abstracting from concurrent changes in marginal costs, my model predicts that markups faced by poor consumers increase relative to those faced by the rich during recessions. In the data, where price changes are also driven by marginal cost changes, relative prices for poor consumers remain countercyclical.

To establish this fact, I compute Törnqvist inflation indices separately for poor and rich households. That is, for $h \in \{\text{poor, rich}\}$, I define

$$\text{inflation}_{h,t} \equiv \exp\left(\sum_{i \in \mathscr{I}} \frac{\text{share}_{i,h,t} + \text{share}_{i,h,t-1}}{2} \log\left(\frac{\text{price}_{i,t}}{\text{price}_{i,t-1}}\right)\right) \tag{28}$$

where $\text{share}_{i,h,t}$ is the average expenditure share of consumer segment h on barcode i. Crucially, in equation (28), barcode-level inflation rates are averaged across households and do not differ across consumer segments. Consequently, any disparity in $\text{inflation}_{h,t}$ stems from differences in consumption choices rather

Figure 7: Inflation for Poor Households is Higher in Recessions



This graph illustrates differences in Törnqvist inflation indices for poor vs rich households. The left-hand side panel plots a time series of the inflation gap $\text{inflation}_{\text{poor},t} - \text{inflation}_{\text{rich},t}$ with the distinction between poor and rich based on income levels. The right-hand side panel plots a time series for the same inflation gap distinguishing households based on expenditures. By construction, inflation indices only reflect product choice and abstract from search behavior.

than search behavior. Demarcating poor versus rich consumers, I present results for two different classification schemes. First, I identify as rich those households whose income exceeds their (region- and time-specific) median. Alternatively, I categorize as rich those consumers with above-median total spending on NielsenIQ barcodes.[9] The inflation gap between poor and rich consumers is defined as

$$\Delta \text{inflation}_{\text{consumers},t} \equiv \text{inflation}_{\text{poor},t} - \text{inflation}_{\text{rich},t}. \tag{29}$$

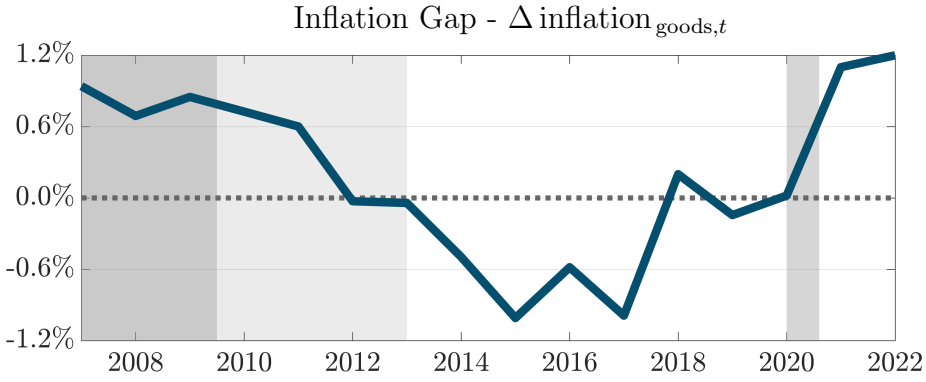
Figure 7 plots the time series of this inflation gap. Consistent with my model predictions, prices faced by poor consumers increase relative to those faced by the rich – both during and in the aftermath of recessions.[10]

In my model, this pattern emerges because cheaper, low-quality options become

[9] Using income rather than expenditure as a proxy for economic status addresses the lack of data on substitutes outside of NielsenIQ (e.g. grocery vs. restaurant spending; a high grocery bill doesn't necessarily imply affluence). Focusing on expenditures is also reasonable, as NielsenIQ lacks wealth data, and income alone might inadequately depict a household's financial situation.

[10] In that sense, recessions are a pivotal driver of inflation inequality as documented by e.g. Jaravel (2019). In the long run, disparities in inflation rates are due to a somewhat muted reversal in normal times.

Figure 8: Inflation for Cheap Goods is Higher in Recessions



This graph plots a time series of the difference in Törnqvist inflation indices for cheap versus premium goods. The distinction between cheap and premium goods is based on time-averaged barcode-level premium scores.

relatively more expensive during recessions. To confirm that this is also true in the data, I compute Törnqvist indices across the entire population but separately for a partition of NielsenIQ barcodes $\{\mathscr{I}_k\}$ where $k \in \{\text{cheap, premium}\}$. That is,

$$\text{inflation}_{k,t} \equiv \exp\left(\sum_{i \in \mathscr{I}_k} \frac{\text{share}_{i,t} + \text{share}_{i,t-1}}{2} \log\left(\frac{\text{price}_{i,t}}{\text{price}_{i,t-1}}\right)\right). \tag{30}$$

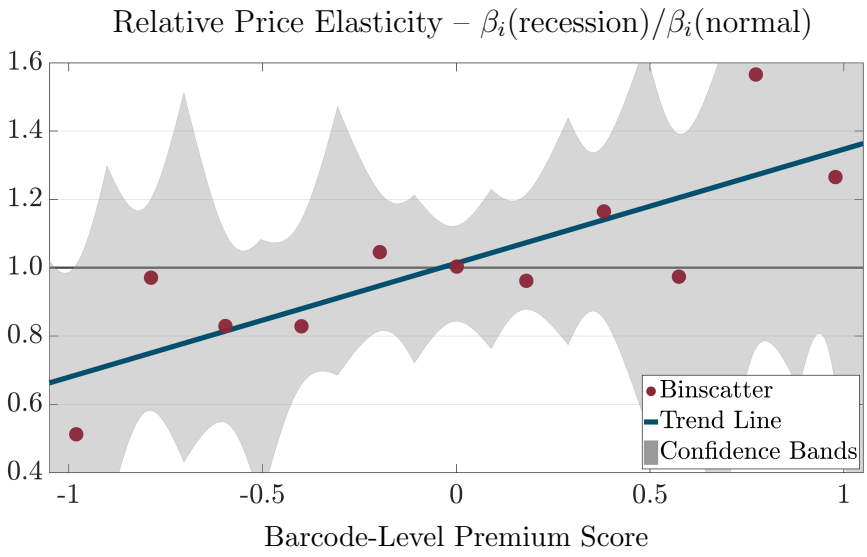
Here, cheap goods are those for which sales-weighted region- and time-average of premium_{irt} is below zero – and vice versa for $\mathscr{I}_{\text{premium}}$. I average premium scores over time to ensure that any correlation with inflation rates is not a mechanical artifact.[11] The inflation gap between cheap and expensive varieties is

$$\Delta \text{inflation}_{\text{goods},t} \equiv \text{inflation}_{\text{cheap},t} - \text{inflation}_{\text{premium},t}. \tag{31}$$

Figure 8 plots the time series of this inflation gap. Consistent with my model predictions, the relative price of cheaper varieties increases during recessions and in their aftermath. Cavallo and Kryvtsov (2024) document this phenomenon of *cheapflation* across many countries following the Covid-19 pandemic.

[11] Specifically, this precludes a scenario where higher inflation on cheaper goods is due to, say, a mean-reverting component in marginal costs.

Figure 9: The Rich Become Less Elastic Toward Cheap Goods in Recessions

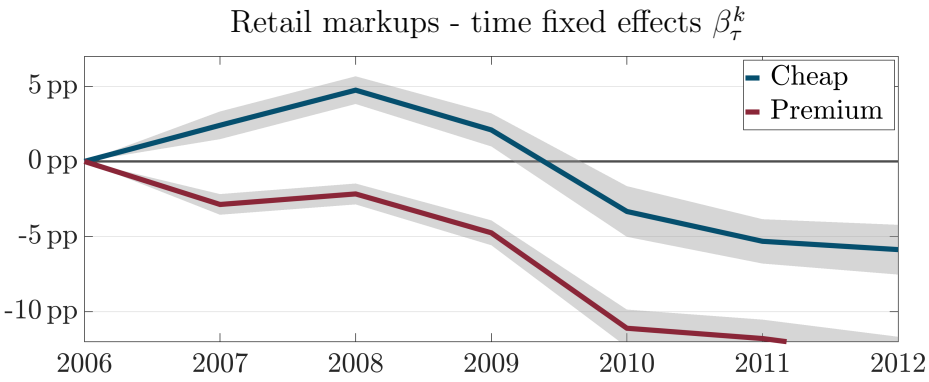


This graph depicts a binscatter of barcode-level premium scores on the x-axis against the corresponding relative price elasticities among wealthier households during and outside of the Great Recession.

Price elasticities during the Great Recession. Another key prediction of my model is that wealthy households become less responsive to price changes for cheaper goods as they shift their spending toward these items during recessions. To check this prediction with data, I estimate barcode-level price elasticities among wealthier consumers both during and outside the Great Recession. For each barcode, I then calculate the ratio of price elasticities during the recession compared to those outside it. Whenever this ratio is less than one for a particular barcode i, it indicates that wealthy households became less responsive to price changes for that specific barcode during the Great Recession. Figure 9 displays a binscatter plot of barcode-level premium scores on the x-axis against this ratio of price elasticities on the y-axis. The graphs reveals that, during the Great Recession, affluent households became less price-elastic toward cheaper goods while exhibiting increased price elasticity toward more expensive items.

Retail markups during the Great Recession. My model predicts that as a recession hits, markups on cheaper, low-quality goods increase, whereas those on

Figure 10: Retail Markups on Cheap Products Increase in Recessions



This graph depicts percentage point time fixed effects for retail markups on cheap versus expensive varieties during and in the aftermath of the Great Recession relative to a base period in 2006.

more expensive, higher-quality goods decline. To explore this prediction in the context of data on retail markups, I use PriceTrak PromoData to gather regional wholesale cost information. Combined with price data from NielsenIQ, I calculate region- and time-specific average retail markups. That is,

$$\text{markup}_{irt} = \frac{\text{price}_{irt}}{\text{wholesale-cost}_{irt}} \tag{32}$$

where price_{irt} is the quantity-weighted average price of barcode i in region r at time t and $\text{wholesale-cost}_{irt}$ are the corresponding average "marginal" costs incurred by retailers. Based on time averages of my premium scores, I then partition the universe of NielsenIQ barcodes $\{\mathscr{I}_k\}$ where $k \in \{\text{cheap, premium}\}$. For each partition, I run the regression

$$\text{markup}_{irt} = \alpha_i^k + \alpha_r^k + \textstyle\sum_{\tau=2007}^{2012} \beta_\tau^k \times \mathbb{1}\{t = \tau\} + \epsilon_{irt}^k \tag{33}$$

including and observations (i, r, t) iff $i \in \mathscr{I}_k$. Figure 10 plots the time series of β_τ^k for cheap versus expensive barcodes. Relative to the (pre-crisis) base period in 2006, retail markups for cheaper, low-quality goods increased immediately following the onset of the Great Recession. Conversely, markups for higher-quality

goods decreased. In the recession's aftermath, retail markups dropped across all products, but this decline was markedly more pronounced among expensive items. The initial divergence in markups aligns perfectly with my model predictions. It is important to note that the determination of retail markups, unlike production markups in my model, is also influenced by unmodeled factors related to across-store competition and within-store cannibalization.

6 Redistribution in General Equilibrium

Since the markup channel amplifies real consumption inequality during recessions, policymakers concerned about inequality might naturally consider implementing redistributive stabilization policies. In this section, I explore the impact of such policy on product market competition within a Bewley-Aiyagari-Hugget model. I find that redistributive stabilization policies worsen the effects of the markup channel.

6.1 General equilibrium model

Since redistribution to the poor typically distorts incentives to save or work, policymakers face a trade-off between equity and efficiency. Consequently, assessing redistributive policies warrants a general equilibrium model. In this section, I embed the markup channel into a Bewley-Aiyagari model with elastic labor supply.

Household behavior. Time is discrete, and the economy is populated by a continuum of infinitely-lived households. Each household consumes a bundle of different-quality varieties. The intratemporal consumption allocation decision is governed by the nonhomothetic preferences from equations (7) and (8). For tractability, I maintain that $|\mathscr{S}| = \mathfrak{c}$ but posit that sector compositions are perfectly symmetric. Households differ in their labor market ability e and choose how many hours of labor to supply. Labor market ability follows a Markov chain

$e' \sim H(e' \,|\, e)$. To insure against idiosyncratic income risk, households save in a single safe asset.

Households enter a period with endogenously chosen, yet pre-determined, asset holdings a as well as an exogenous draw of their idiosyncratic labor market ability e. Written recursively, the consumers' problem is to choose $\big(c, h, a'\big)$ in order to solve

$$V(a, e) = \max \left\{ u(c, h) + \beta \, \mathbb{E}\Big[V(a', e') \,\Big|\, e \Big] \right\} \tag{34}$$

where optimization is subject to a budget constraint

$$p\big(c, \boldsymbol{p}\big)\, c + \bar{p}\, a' = \big(1 + r\big)\, \bar{p}\, a + (1 - \tau)\, w\, e\, h + \pi(a) + T \tag{35}$$

as well as no-borrowing condition $a' \geq 0$. In the background there is a perfectly competitive investment sector that combines varieties in equal proportions to produce an investment good with relative price $\bar{p}$. Firm ownership is proportional to a strictly increasing function of asset holdings. That is, aggregate profits π are are distributed according to

$$\pi(a) = \frac{f(a)}{\int f(a)\, \Gamma(da, de)}\, \pi \tag{36}$$

where Γ is the endogenous stationary distribution over idiosyncratic consumer states. The key departure from a standard incomplete market setting with elastic labor supply is that real consumption enters the budget constraint nonlinearly. Specifically, the nonhomothetic ideal price index in equation (35) is given as

$$p(c, \boldsymbol{p}) = \left(\sum_{q=1}^{Q} \sum_{i=1}^{N} \varphi_q \, p_{iq}^{1-\sigma} c^{(1-\sigma)(\xi_q - 1)} \right)^{\frac{1}{1-\sigma}} . \tag{37}$$

This object succinctly encapsulates all the intricacies of intratemporal consumption allocation under the nonhomothetic preferences discussed in Section 2.

The households' problem yields policy functions for real consumption $c(a, e)$, hours worked $h(a, e)$, as well as savings $a'(a, e)$. The implied policy function for

nominal spending is

$$y(a,e) = p\big(c(a,e), \boldsymbol{p}\big)\, c(a,e). \tag{38}$$

Since spending is now a choice variable, the expenditure distribution is endogenous and prescribed by the pushforward measure $G = \Gamma \circ y^{-1}$. This endogenous expenditure distribution is crucial for product market competition.

To characterize how households trade off spending, leisure, and savings, I next present the first-order conditions that shape consumer decision making. To that end, I first define the marginal price of real consumption as the partial derivative of nominal spending with respect to real consumption.[12]

$$\tilde{p}(c, \boldsymbol{p}) \;\equiv\; p(c, \boldsymbol{p}) + \frac{\partial\, p(c, \boldsymbol{p})}{\partial\, c}\, c. \tag{39}$$

The consumers' intratemporal consumption leisure trade-off is then governed by

$$u_h(c,h) = -u_c(c,h)\, \frac{(1-\tau)\, w\, e}{\tilde{p}\big(c, \boldsymbol{p}\big)} \tag{40}$$

while the intertemporal consumption savings trade-off is given through

$$\mathbb{E}\left[\, \beta\, \frac{u_c(c', h')}{u_c(c,h)}\, \frac{\tilde{p}\big(c, \boldsymbol{p}\big)}{\tilde{p}\big(c', \boldsymbol{p}'\big)}\, \big(1+r\big) \,\middle|\, e \,\right] \leq 1. \tag{41}$$

Firm behavior. Firms maximize profits vis-à-vis the now-endogenous expenditure distribution G. Specifically, they operate a constant returns to scale production technology such that

$$c_{iq} = z_q\, \ell_{iq}^{1-\alpha} k_{iq}^{\alpha}. \tag{42}$$

Cost minimization, therefore, dictates that firms produce under constant marginal cost with

$$\lambda_q = \frac{1}{z_q} \left(\frac{w}{1-\alpha}\right)^{1-\alpha} \left(\frac{r\,\bar{p}}{\alpha}\right)^{\alpha}. \tag{43}$$

The price setting protocol is precisely as beforehand. Given consumer policies,

[12] For a natural cardinalization of homothetic preferences with $\xi_q = 1$ for all q, the marginal price of real consumption is just the familiar constant homothetic CES price index.

profit-maximizing prices are pinned down through a Lerner-type expression for markups

$$\mu_q(\boldsymbol{p}, g) \equiv \frac{p_q}{\lambda_q} = \frac{\int \varepsilon_q(y, \boldsymbol{p})\, \tilde{c}_q(y, \boldsymbol{p}, g)\, g(y)\, dy}{\int \varepsilon_q(y, \boldsymbol{p})\, \tilde{c}_q(y, \boldsymbol{p}, g)\, g(y)\, dy - 1}. \tag{44}$$

where $g(y)$ satisfies $\Gamma \circ y^{-1}(da, de) = g(y)\, dy$.

Equilibrium. The stationary equilibrium in this economy is characterized by a vector $\big(r, w, \boldsymbol{p}, \pi, T\big)$ such that households optimize according to (34) and (35) and different-quality prices are consistent with $p_q = \mu_q\big(\boldsymbol{p}, g\big)\, \lambda_q$ for all $q \in \big\{1, \ldots, Q\big\}$. Moreover, market clear such that

$$r\, \bar{p} \int a\, \Gamma(da, de) = \alpha \sum_{q=1}^{Q} \sum_{i=1}^{N} \lambda_q \int c_q(a, e)\, \Gamma(da, de) \tag{45}$$

$$w \int e\, h(a, e)\, \Gamma(da, de) = (1 - \alpha) \sum_{q=1}^{Q} \sum_{i=1}^{N} \lambda_q \int c_q(a, e)\, \Gamma(da, de) \tag{46}$$

and the government runs a balanced budget with

$$\tau\, w \int e\, h(a, e)\, \Gamma(da, de) = T. \tag{47}$$

The stationary distribution satisfies

$$\Gamma\big(\mathscr{A} \times \mathscr{E}\big) = \int \mathbb{1}\big\{a'(a, e) \in \mathscr{A}\big\} \sum_{e' \in \mathscr{E}} H(e'|e)\ \Gamma(da, ae) \tag{48}$$

for all Borel-sets $\mathscr{A}$ and $\mathscr{E}$ and $g(y)$ is the Radon-Nikodym derivative of the pushforward $\Gamma \circ y^{-1}$. The price of investment is simply given as $\bar{p} = \sum_q p_q / Q$.

Calibration. I calibrate my general equilibrium model at the quarterly frequency to match data moments of the joint income and wealth distribution. The parameterization of the nonhomothetic ideal price index follows the same strategy I adopted in Section 4. Parameter values, calibration targets, and model fit are reported in Table 7.

Table 7: Dynamic GE Model – Calibration and Model Fit

Parameter	Value	Significance	Target	Data	Model
α	0.33	Capital elasticity of output	Assigned	-	-
γ	2	Inverse Frisch elasticity	Assigned	-	-
β	0.9572	Discount rate	Average wealth to income	16.4	16.9
θ	3.46	Constant relative risk aversion	Top 10% wealth share	0.49	0.46
μ	1.36	Mean labor market ability	Gini income	0.39	0.42
s	0.045	Dispersion labor market ability	Top 10% income share	0.31	0.32
ρ	0.968	Persistence labor market ability	Persistence income	0.98	0.97
η	1.55	Across-sector substitution	Aggregate markup	1.43	1.40
σ	12	Within-sector substitution	Sales HHI	0.22	0.25
$\varphi_{\text{high}}/\varphi_{\text{low}}$	1.22	Demand shifters	Polarization (mid/poor)	6.18	5.98
$\xi_{\text{high}}/\xi_{\text{low}}$	0.523	Nonhomotheticities	Premium index (rich/poor)	1.21	1.17
$z_{\text{high}}/z_{\text{low}}$	0.84	Relative productivity	Relative price	1.24	1.22

This table reports the jointly calibrated parameters of my dynamic general equilibrium model. Specifically, I match various moments of the joint wealth and income distribution in PSID data. My calibration target for the aggregate markup is aligned with the model-implied markup from my partial equilibrium exercise. My calibration strategy for the parameters speaking to quality distinctions in production and consumption patterns follows my approach from Section 4.

6.2 The markup channel and redistributive policy

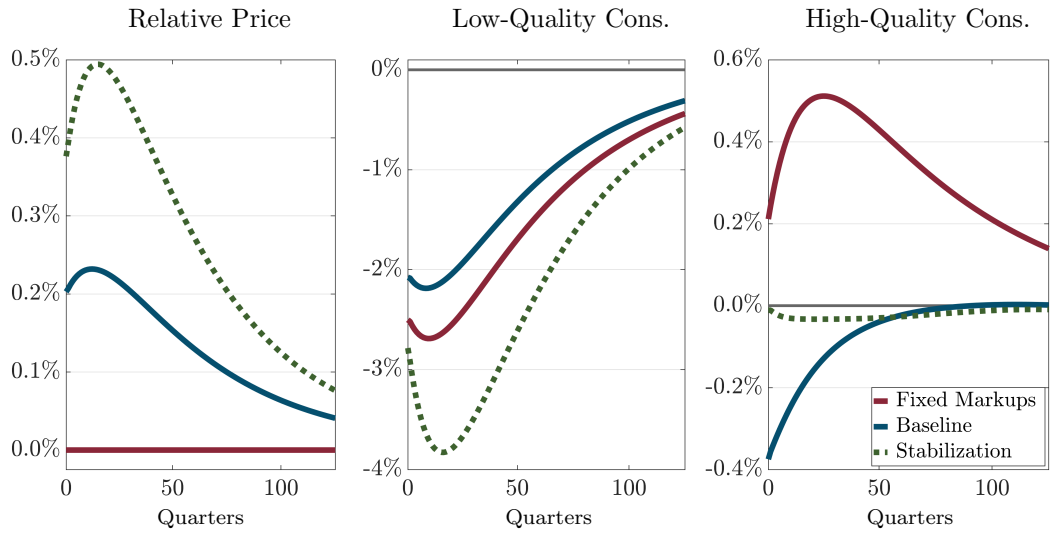
In this section, I examine how a recession affects relative prices and the cyclicality of different-quality outputs within this framework. I find that redistributive policy further increases the relative price of low-quality goods.

The markup channel in general equilibrium. I engineer a recession through a persistent aggregate TFP shock Θ_t that symmetrically decreases productivity across all quality tiers. That is, for each $q \in \{\text{low, high}\}$, productivity is given by

$$z_{q,t} = z_q \exp(\Theta_t).$$

For illustration, $\{\Theta_t\}$ encodes a -5% MIT shock hitting the economy at time $t = 0$ and decaying exponentially with a persistence of 0.95. In order to isolate the effects of the markup channel, Figure 11 shows the time paths of relative prices and

Figure 11: Time Path of Aggregates in Response to MIT Shock to Θ



This Figure depicts the time path of the relative price of low-quality goods, aggregate consumption of low-quality goods, and aggregate consumption of high-quality goods in response to a -5% TFP shock. The stabilization parameter is $\psi = -2$.

consumption in two distinct scenarios: without and with markup adjustments.[13]

First, I counterfactually prevent firms from reoptimizing markups by keeping $\mu_{q,t}$ fixed at its pre-recession value. Since marginal cost changes are identical across producers, this suppression of the markup channel eliminates all changes in the relative price of different-quality goods. As the recession hits, households shift from more expensive, higher-quality goods to more affordable options. The red lines in the middle and right panels of Figure 11 illustrate that, without markup adjustments, high-quality consumption significantly declines, whereas consumption of lower-quality goods increases relative to its steady-state level.

Second, I allow producers to adjust their markups to maximize profits. As households shift toward lower-quality consumption, lower-quality producers gain market share and charge higher markups; the relative price of low-quality goods increases during the recession. This rise in relative price, in turn, moderates the shift toward low-quality varieties; so much so that aggregate consumption of low-

[13]The computation of these time-paths follows the first-order approximation from Bhandari, Bourany, Evans, and Golosov (2023).

quality goods actually falls. The markup channel, therefore, dampens the cyclical fluctuations of high-quality consumption while amplifying those of low-quality consumption.

Redistribution through automatic stabilization. To assess the interaction between the markup channel and redistributive stabilization policy, I study the implementation of an automatic stabilizer $\psi\,\Theta_t$.[14] By setting the policy parameter $\psi < 0$, a policymaker imposes additional taxes on labor income and redistributes the corresponding revenue through lump-sum rebates. The resulting budget constraint is given by

$$p\big(c_t, \boldsymbol{p}_t\big)\, c_t + \bar{p}_t\, a_{t+1} = \big(1 + r_t\big)\, \bar{p}_t\, a_t + \big(1 - \tau - \psi\,\Theta_t\big)\, w_t\, e_t\, h_t + \pi(a_t) + T_t + S_t \quad (49)$$

where the lump-sum transfer S_t satisfies

$$S_t = \psi\ \Theta_t \int h_t(a, e)\, e\ \Gamma_t(da, de). \quad (50)$$

Note that this policy intervention constitutes a redistribution toward households with lower labor income. Evaluating the impact of this automatic stabilizer requires a general equilibrium framework, as an increase in labor taxes is bound to distort incentives for labor supply.

The dotted green line in the leftmost panel of Figure 11 depicts the time path of the relative price of low-quality goods for $\psi = -2$, meaning that a 1% drop in aggregate TFP results in a 2 percentage point increase in labor taxes. The relative price increase is more than twice as large as in the scenario without stabilization. Intuitively, when redistributing to the poor, the policymaker redirects funds that would have been spent on high-quality goods toward spending on lower-quality options. Consequently, lower-quality producers gain even more market share and charge even higher markups. While the markup channel itself

[14] While optimal policy design is beyond the scope of this paper, the presence of curvature in my model economy implies that, in the wake of a recession, a policy authority that is concerned with inequality finds it desirable to redistribute towards the poor.

increases real consumption inequality, its interaction with redistributive policy interventions generates even more *cheapflation*.

7 Conclusion

This paper shows that the well-documented shift in demand towards lower-quality goods during recessions leads to higher prices for poorer consumers. The primary mechanism behind this phenomenon is the *markup channel*: as consumers switch to lower-quality goods, producers of these items gain market power and raise their markups. This mechanism highlights recessions as a significant driver of *cheapflation*.

Given that the markup channel exacerbates real consumption inequality, policymakers might consider redistributing resources toward the poor during recessions. By incorporating the markup channel into a Bewley-Aiyagari-Hugget framework, this paper demonstrates that such redistribution can actually increase the relative price of lower-quality goods even further. Since the negative impact on real consumption inequality arises from a disruption in product market competition, the insights from this paper advocate for product market interventions. Future research could explore quality-specific subsidies, such as those implemented through food stamps, as a potential solution.

References

Aguiar, Mark and Mark Bils (2015). "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review*, 105, 2725-2756.

Aguiar, Mark and Erik Hurst (2005). "Consumption versus Expenditure." *Journal of Political Economy*, 113, 919-948.

Albrecht, James, Guido Menzio, and Susan Vroman (2023). "Vertical Differentiation in Frictional Product Markets." *Journal of Political Economy: Macroeconomics*, 1, 586-632.

Anderson, Eric, Sérgio Rebelo, and Arlene Wong (2020). "Markups Across Space and Time."

Argente, David and Munseob Lee (2021). "Cost of Living Inequality during the Great Recession." *Journal of the European Economic Association*, 19, 913-952.

Atkeson, Andrew and Ariel Burstein (2008), "Pricing-to-Market,Trade Costs, and International Relative Prices." *American Economic Review*, 98,1998-2031.

Becker, Jonathan, Chris Edmond, Virgiliu Midrigan, and Daniel Y. Xu (2024). "National Concentration, Local Concentration, and the Spatial Distribution of Markups."

Benkard, C. Lanier, Ali Yurukoglu, and Anthony Lee Zhang (2021). "Concentration in Product Markets." *National Bureau of Economic Research.*

Berry Steven, James Levinsohn, and Ariel Pakes (1995). "Automobile Prices in Market Equilibrium." *Econometrica*, 63.

Bils, Mark and Peter J. Klenow (2001). "Quantifying Quality Growth." *American Economic Review*, 91, 1006-1030.

Bisgaard-Larsen, Rasmus and Christoffer J. Weissert (2020). "Quality and Consumption Basket Heterogeneity."

Boar, Corina, and Virgiliu Midrigan (2023). "Markups and Inequality." *National Bureau of Economic Research.*

Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch (2021). "Some Unpleasant Markup Arithmetic: Production Function Elasticities and their Estimation from Production Data." *Journal of Monetary Economics*, 121, 1-14.

Boppart, Timo (2014). "Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences." *Econometrica*, 82, 2167-2196.

Bornstein, Gideon and Alessandra Peter (2024), "IneQuality."

Buera, Fransisco J. and Joseph P. Kaboski (2009). "Can Traditional Theories of Structural Change Fit the Data?" *Journal of the European Economic Association*, 7, 469-477.

Burstein, Ariel, Martin Eichenbaum, and Sérgio Rebelo (2005). "Large Devaluations and the Real Exchange Rate." *Journal of Political Economy*, 113, 742-784.

Burstein, Ariel, Vasco M. Carvalho, and Basile Grassi (2020). "Bottom-up Markup Fluctuations." *National Bureau of Economic Research.*

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff (2014). "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood." *American Economic Review*, 104, 2633-2679.

Cirelli, Fernando (2022). "Bank-Dependent Households and the Unequal Costs of Inflation."

Comin, Diego, Danial Lashkari, and Martí Mestieri (2021). "Structural Change with Long-Run Income and Price Effects." *Econometrica*, 89, 311-374.

Deaton, Angus, and John Muellbauer (1980). "Economics and Consumer Behavior." *Cambridge University Press.*

DellaVigna, Stefano and Matthew Gentzkow (2019). "Uniform Pricing in US Retail Chains." *The Quarterly Journal of Economics*, 134, 2011-2084.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu (2013). "How Costly are Markups?" *Journal of Political Economy* 131, 000-000.

Fajgelbaum, Pablo D., Gene M. Grossman, and Elhanan Helpman (2011). "Income Distribution, Product Quality, and International Trade." *Journal of Political Economy*, 119, 721-765.

Ferraro, Domenico and Vytautas Valaitis (2022). "Wealth and Hours."

Handbury, Jessie (2021). "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living across US Cities." *Econometrica* 89, 2679-2715

Hanoch, Giora (1975). "Production and Demand Models With Direct or Indirect Implicit Additivity." *Econometrica*, 43, 395-419.

Hausman, Jerry A. (1996). "Valuation of New Goods under Perfect and Imperfect Competition." *The Economics of New Goods*.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante (2020). "The Rise of US Earnings Inequality: Does the Cycle drive the Trend?" *Review of Economic Dynamics*, 37, 181-204.

Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi (2014). "Growth and Structural Transformation." *Handbook of Economic Growth*, 2, 855-941.

Hitsch, Günter J., Ali Hortacsu, and Xiliang Lin (2019). "Prices and Promotions in US Retail Markets: Evidence from Big Data." *NBER Working Paper*.

Jaimovich, Nir, Sérgio Rebelo, and Arlene Wong (2019). "Trading down and the Business Cycle." *Journal of Monetary Economics*, 102, 96-12.

Jaravel, Xavier (2019). "The Unequal Gains from Product Innovations: Evidence from the US Retail Sector." *The Quarterly Journal of Economics*, 134, 715-783.

Jaravel, Xavier, and Danial Lashkari (2024). "Measuring Growth in Consumer Welfare with Income-Dependent Preferences: Nonparametric Methods and Estimates for the United States." *The Quarterly Journal of Economics*, 139, 477-532.

Jaravel, Xavier and Alan Olivi (2021). "Prices, Non-homotheticities, and Optimal Taxation."

Jørgensen, Casper N. and Leslie Shen (2022), "Consumption Quality and the Welfare Implications of Business Cycle Fluctuations."

Kaplan, Greg and Guido Menzio (2015). "The Morphology of Price Dispersion." *International Economic Review* 56, 1165-1206.

Kaplan, Greg and Guido Menzio (2016)."Shopping Externalities and Self-Fulfilling Unemployment Fluctuations." *Journal of Political Economy* 1243, 771-825.

Kongsamut, Piyabha, Sérgio Rebelo, and Danyang Xie (2001). "Beyond Balanced Growth." *Review of Economic Studies*, 68, 869-882.

Marto, Ricardo (2023). "Structural Change and the Rise in Markups."

Matsuyama, Kiminori (2023). "Non-CES Aggregators: A Guided Tour." *Annual Review of Economics*, 15, 235-265.

Mongey, Simon and Make Waugh (2024), "Pricing Inequality."

Nord, Lukas (2022). "Shopping, Demand Composition, and Equilibrium Prices." *SSRN*.

Sangani, Kunal (2023). "Markups Across the Income Distribution: Measurement and Implications."

Straub, Ludwig (2019). "Consumption, Savings, and the Distribution of Permanent Income."

Wachter, Jessica A. and Motohiro Yogo (2010). "Why Do Household Portfolio Shares Rise in Wealth?" *The Review of Financial Studies*, 23, 3929-3965.

Appendix A - Additional Results & Figures

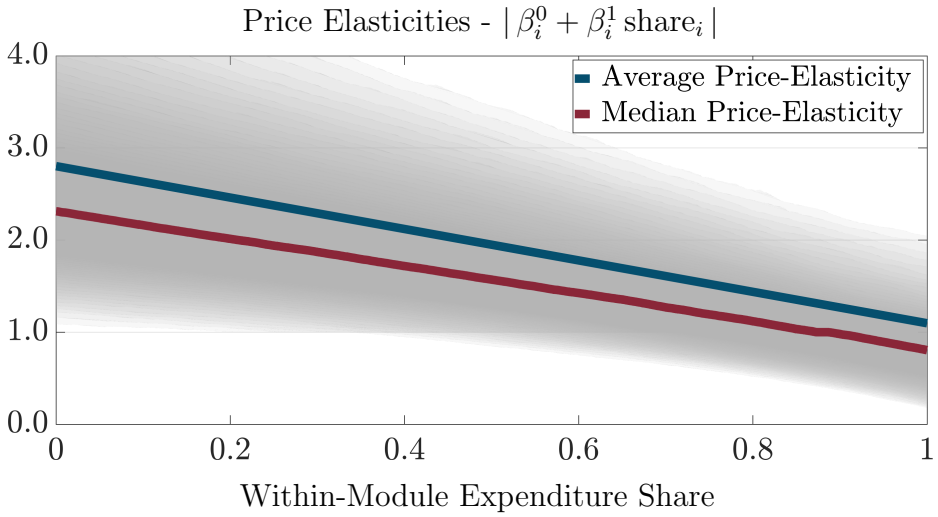
Fact 3 via stratification of population. To establish fact 3 from Section 3 in the data, here, I estimate barcode-specific price elasticities as a function of within-module expenditure shares. Specifically, I conduct barcode-level instrumental variable regressions to estimate a log-linearization of my demand system from (16). My approach leverages the fact that model-implied price elasticities depend on household characteristics only through expenditure shares. In particular, I interact household h's expenditure share for variety i within product module $m(i)$ with household-specific log prices. The resulting regression equation is:

$$\begin{aligned}\log \text{quantity}_{iht} = \; & \alpha_{ih} + \alpha_{it} + \alpha_{ir(h)} + \left(\beta_i^0 + \beta_i^1 \, \text{share}_{iht}\right) \log \text{price}_{iht} \\ & + \sum\nolimits_{j \in \mathscr{C}_{iht}} \beta_{ij} \log \text{price}_{jr(h)t} + \gamma_i \log \text{expenditure}_{ht} + \epsilon_{iht}.\end{aligned} \tag{51}$$

In this regression, $\beta_i(x) \equiv \left| \, \beta_i^0 + \beta_i^1 \, x \, \right|$ can be interpreted as barcode i's own price elasticity among consumers who allocate an expenditure shares x on barcode i in product module $m(i)$. The regression controls for both household expenditures and a carefully constructed set of household-specific competitors for each barcode i. Using data on shopping trips from the consumer panel and store-level pricing information from NielsenIQ's retail-scanner data, I ensure that $\mathscr{C}_{iht}$ is comprised of barcodes $j \in m(i)$ that are actually available to household h at $\text{price}_{jr(h)t}$.

In order to address potential endogeneity issues in the relationship between quantity_{iht} and price_{iht}, note that idiosyncratic determinants of the quantity choice of a particular household are likely orthogonal to retail prices. Therefore, the reverse causality concern boils down to the presence of local demand shocks that are observable to retailers (and thus reflected in pricing) but unobservable to the econometrician. To address this concern, I construct a set price instruments in the spirit of Hausman (1996). Specifically, I instrument price_{iht} with the economy-wide average price for barcode i in year t, excluding observations from region $r(h)$. Moreover, since consumption choices along the premium margin are highly correlated with income, I instrument expenditure shares with household income_{ht}. Price elasticities are, therefore, identified by within-household/region/time price variation that is explained by economy-wide price movements. The identifying assumption is that pricing decisions that apply to barcodes throughout the nation are orthog-

Figure 12: Cross-sectional Distribution of Price Elasticities



This graph depicts the distribution of $i \mapsto \beta_i(x)$ across the universe of NielsenIQ barcodes. The graph shows that both average and median price elasticities decrease as within-module spending shares x increase. The shaded area extends from the 5th to 95th percentile of the cross-sectional distribution.

onal to local demand conditions. This exclusion restriction is broadly consistent with evidence of uniform pricing documented by DellaVigna and Gentzkow (2019). Since all coefficients in (51) are indexed at the barcode level, under mild clustering conditions on ϵ_{iht}, all regressions are run barcode by barcode.

Figure 12 illustrates the distribution of $i \mapsto \beta_i(x)$ across the universe of NielsenIQ barcodes. The graph shows that both average and median price elasticities decrease as within-module spending shares x increase. That is, as consumers increase their relative spending on a particular variety, they become less price-elastic with respect to it. Consequently, from Figure 12, we can deduce that poor consumers, who oftentimes concentrate their within-sector spending on the least expensive options available, inelastically choose these low-cost alternatives. Similarly, wealthier consumers are inelastic regarding premium goods, as they devote a significant portion of their spending to these pricier varieties. By contrast, middle-class consumers, whose spending is more evenly distributed across a range of varieties along the premium margin, display greater price elasticity in either direction.

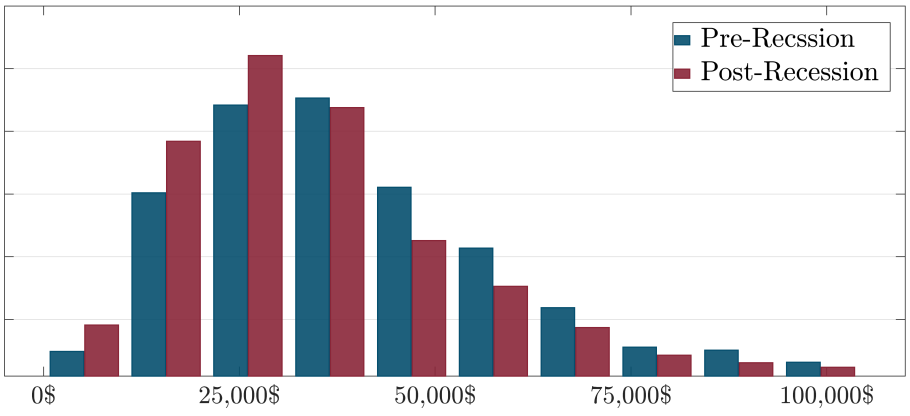
When linking this empirical observation to model-implied price elasticities, note that in a homothetic model environment, expenditure shares are uniform across the entire population. Consequently, there is no cross-sectional heterogeneity in price elasticities. However, as shown in Figure 1, nonhomothetic preferences lead

to variation in expenditure shares on different-quality products across the expenditure distribution. Figure 2 further illustrates that model-implied price elasticities generally decrease as expenditure shares increase.

A remark is in order: While the price elasticities in Figure 12 align with other estimates based on NielsenIQ data, such as those from Hitsch, Hortacsu, and Lin (2019), they are significantly lower than standard macro estimates of elasticities of substitution. One possible explanation for this discrepancy is that elasticities of substitution within grocery stores are, in fact, relatively low, and that retail markups are largely influenced by competition among stores. The key insight from this exercise is qualitative: when facing a fixed product assortment, households become less price elastic for a particular variety as their expenditure share on it increases.

Histogram of the expenditure distribution pre- and post-crisis.

Figure 13: Expenditure Distribution - 2006 vs 2012



This figure is a histogram of the symmetrically PCE-deflated expenditures in PSID data for 2006 as well as 2012. During this time period, i.e., during the Great Recession, average spending decreased by 15.9% while the 75/25 percentile ratio of the distribution decreased from 3.02 to 2.89.

Appendix B - Model and Derivations

Households' program. Consumers take as given the economy's price vector $\boldsymbol{p} = \left(\boldsymbol{p}_s : s \in \mathscr{S}\right)$ where $\boldsymbol{p}_s = \left(p_{iqs} : q \in \mathscr{Q} \text{ and } i \in \mathscr{I}_{qs}\right) \forall\ s \in \mathscr{S}$. They choose allocations $\left\{c_{iqs}\right\}$ to minimize the expenditure necessary to attain real consumption c. That is, they solve

$$\inf_{\{c_{iqs},\, c_s\}} \int_{\mathscr{S}} \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} p_{iqs} c_{iqs}\, ds \tag{52}$$

subject to

$$\int_{\mathscr{S}} \left(\frac{c_s}{c}\right)^{\frac{\eta-1}{\eta}} ds = 1 \tag{53}$$

and

$$\sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} \left(\frac{\varphi_q}{c_s^{(1-\sigma)(1-\xi_q)}}\right)^{\frac{1}{\sigma}} \left(\frac{c_{iqs}}{c_s}\right)^{\frac{\sigma-1}{\sigma}} = 1 \quad \forall \quad s \in \mathscr{S}. \tag{54}$$

Intermediate Hicksian demand. Invoking a standard separation theorem, the first step is to minimize within-sector expenditures to attain a given level of sectoral consumption c_s. To that end, households solve

$$\min_{\{c_{iqs}\}} \left\{ \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} p_{iqs} c_{iqs} \;\middle|\; \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} \varphi_q^{\frac{1}{\sigma}} \left(\frac{c_{iqs}}{c_s^{\xi_q}}\right)^{\frac{\sigma-1}{\sigma}} = 1 \right\}. \tag{55}$$

With λ being the Lagrange-multiplier on the nonhomothetic CES aggregator, the first-order conditions with respect to c_{iqs} are

$$p_{iqs} = \lambda \frac{\sigma-1}{\sigma} \varphi_q^{\frac{1}{\sigma}} \left(\frac{c_{iqs}}{c_s^{\xi_q}}\right)^{\frac{\sigma-1}{\sigma}-1} \frac{1}{c_s^{\xi_q}}. \tag{56}$$

Multiplying by c_{iqs} and summing over producers and quality-bins we obtain sectoral spending y_s as

$$y_s \equiv \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} p_{iqs} c_{iqs} = \lambda \frac{\sigma-1}{\sigma} \tag{57}$$

where the latter equality follows from the definition of the nonhomothetic CES aggregator in (54). Sectoral expenditure shares are, thus, given as

$$x_{iqs} \equiv \frac{p_{iqs} c_{iqs}(c_s, \boldsymbol{p}_s)}{\sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} p_{iqs} c_{iqs}} = \varphi_q^{\frac{1}{\sigma}} \left(\frac{c_{iqs}}{c_s^{\xi_q}} \right)^{\frac{\sigma-1}{\sigma}}. \tag{58}$$

At this point, we can define nonhomothetic ideal price-index $p_s(c_s, \boldsymbol{p}_s)$ to satisfy

$$y_s = p_s(c_s, \boldsymbol{p}_s)\, c_s = \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} p_{iqs} c_{iqs}. \tag{59}$$

It is instructive to use (57) and (59) and rearrange terms in (56) to obtain

$$c_{iqs} = \varphi_q \left(\frac{p_{iqs}}{p_s(c_s, \boldsymbol{p}_s)} \right)^{-\sigma} c_s^{\sigma + (1-\sigma)\, \xi_q}. \tag{60}$$

Defining the nonhomothetic taste-shifter $\psi_q(c_s)$ as

$$\psi_q(c_s) \equiv \varphi_q\, c_s^{(1-\sigma)(\xi_q - 1)} \tag{61}$$

intermediate Hicksian demand functions are, in turn, concisely written as

$$c_{iqs}(c_s, \boldsymbol{p}_s) = \psi_q(c_s) \left(\frac{p_{iqs}}{p_s(c_s, \boldsymbol{p}_s)} \right)^{-\sigma} c_s. \tag{62}$$

To obtain an expression for $p_s(c_s, \boldsymbol{p}_s)$, expenditure shares can be rewritten as

$$x_{iqs} = \varphi_q \left(\frac{p_{iqs}}{p_s(c_s, \boldsymbol{p}_s)} \right)^{1-\sigma} c_s^{(1-\sigma)(\xi_q - 1)} \tag{63}$$

Since $\sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} x_{iqs} = 1$, it then follows that

$$p_s(c_s, \boldsymbol{p}_s) = \left(\sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} \varphi_q\, p_{iqs}^{1-\sigma} c_s^{(1-\sigma)(\xi_q - 1)} \right)^{\frac{1}{1-\sigma}}. \tag{64}$$

Intermediate Marshallian demand. Marshallian demand is not available in closed-form. We can, however, recover the sectoral nonhomothetic ideal price-index $p_s(y_s, \boldsymbol{p}_s)$ as a function of sectoral spending y_s in a single fixed-point problem. That is, in a slight abuse of notation,

$$p_s(y_s, \boldsymbol{p}_s) = \text{fix} \left\{ p \mapsto \left(\sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} \varphi_q \, p_{iqs}^{1-\sigma} \left(\frac{y_s}{p} \right)^{(1-\sigma)(\xi_q - 1)} \right)^{\frac{1}{1-\sigma}} \right\}. \tag{65}$$

This object, in turn, pins down intermediate Marshallian demand as

$$c_{iqs}(y_s, \boldsymbol{p}_s) = \varphi_q \left(\frac{p_{iqs}}{p_s(y_s, \boldsymbol{p}_s)} \right)^{-\sigma} \left(\frac{y_s}{p_s(y_s, \boldsymbol{p}_s)} \right)^{\sigma + (1-\sigma)\,\xi_q}. \tag{66}$$

Demand for sectoral aggregates. Conditional on the optimal within-sector allocation of consumption, we can next determine the expenditure-minimizing allocation of sectoral real consumption indices. Note that the nonhomothetic ideal price-index recovered in (64) encapsulates optimality of the consumers' within-sector decision problem. As a consequence, at the outer nest, we can think of the households' program as a homothetic expenditure minimization problem with a non-linear pricing structure. That is,

$$\inf_{\{c_s\}} \left\{ \int_{\mathscr{S}} p_s(c_s, \boldsymbol{p}_s) \, c_s \, ds \;\middle|\; \int_{\mathscr{S}} \left(\frac{c_s}{c} \right)^{\frac{\eta-1}{\eta}} ds = 1 \right\}. \tag{67}$$

With λ being the Lagrange-multiplier on the across-sector CES aggregator, the first-order conditions with respect to c_s are such that

$$\frac{\partial \, p_s(c_s, \boldsymbol{p}_s)}{\partial \, c_s} \, c_s + p_s(c_s, \boldsymbol{p}_s) = \lambda \, \frac{\eta - 1}{\eta} \left(\frac{c_s}{c} \right)^{\frac{\eta-1}{\eta} - 1} \frac{1}{c} \tag{68}$$

To arrive at the first-order condition from (12), first off, we recover the partial

derivative of sectoral prices with respect to real consumption as

$$\frac{\partial\, p_s(c_s, \boldsymbol{p}_s)}{\partial\, c_s} = \frac{\partial}{\partial\, c_s}\left[\left(\sum_{q=1}^{Q}\sum_{i=1}^{N_{qs}} \varphi_q\, p_{iqs}^{1-\sigma}\, c_s^{(1-\sigma)(\xi_q-1)}\right)^{\frac{1}{1-\sigma}}\right] \tag{69}$$

$$= p_s(c_s, \boldsymbol{p}_s)^{\sigma} \sum_{q=1}^{Q}\sum_{i=1}^{N_{qs}} \varphi_q\, p_{iqs}^{1-\sigma}\, c_s^{(1-\sigma)(\xi_q-1)-1}(\xi_q - 1) \tag{70}$$

Using (63) and rewriting the derivative as an elasticity, the above expression simplifies to

$$\frac{\partial\, \log p_s(c_s, \boldsymbol{p}_s)}{\partial\, \log c_s} = \sum_{q=1}^{Q}\sum_{i=1}^{N_{qs}} x_{iqs}(c_s, \boldsymbol{p}_s)\, \xi_q - 1. \tag{71}$$

For notational compactness, I define

$$\bar{\xi}_s(c_s, \boldsymbol{p}_s) = \sum_{q=1}^{Q}\sum_{i=1}^{N_{qs}} x_{iqs}(c_s, \boldsymbol{p}_s)\, \xi_q. \tag{72}$$

Multiplying the first-order condition from (68) by c_s we obtain

$$\left[\frac{\partial\, \log p_s(c_s, \boldsymbol{p}_s)}{\partial\, \log c_s} + 1\right] p_s(c_s, \boldsymbol{p}_s)\, c_s = \lambda\, \frac{\eta - 1}{\eta}\left(\frac{c_s}{c}\right)^{\frac{\eta-1}{\eta}}. \tag{73}$$

It, thus, follows that

$$p_s(c_s, \boldsymbol{p}_s)\, c_s\, \bar{\xi}_s(c_s, \boldsymbol{p}_s) = \lambda\, \frac{\eta - 1}{\eta}\left(\frac{c_s}{c}\right)^{\frac{\eta-1}{\eta}}. \tag{74}$$

We can now integrate over sectors and use the definition of across-sector aggregation from (53) to see that

$$\int_{\mathscr{S}} p_s(c_s, \boldsymbol{p}_s)\, c_s\, \bar{\xi}_s(c_s, \boldsymbol{p}_s)\, ds = \lambda\, \frac{\eta - 1}{\eta}. \tag{75}$$

It then immediately follows that

$$\frac{p_s(c_s, \boldsymbol{p}_s)\, c_s\, \bar{\xi}_s(c_s, \boldsymbol{p}_s)}{\int p_s(c_s, \boldsymbol{p}_s)\, c_s\, \bar{\xi}_s(c_s, \boldsymbol{p}_s)\, ds} = \left(\frac{c_s}{c}\right)^{\frac{\eta-1}{\eta}}. \tag{76}$$

Using the fact that $p_s(c_s, \boldsymbol{p}_s)\, c_s = \sum_q \sum_i p_{iqs}\, c_{iqs}(c_s, \boldsymbol{p}_s)$, this can be rewritten as

$$\frac{\sum_q \sum_i p_{iqs}\, c_{iqs}(c_s, \boldsymbol{p}_s)\, \xi_q}{\int \sum_q \sum_i p_{iqs}\, c_{iqs}(c_s, \boldsymbol{p}_s)\, \xi_q\, ds} = \left(\frac{c_s}{c}\right)^{\frac{\eta-1}{\eta}}. \tag{77}$$

Bertrand competition. To obtain an expression for the Bertrand price-elasticity of variety (k, r, s), I, first off, take the logarithm of the intermediate Hicksian demand function from (62). That is,

$$\log c_{krs}(c_s, \boldsymbol{p}_s) = \log \varphi_r - \sigma \log p_{krs} + \sigma \log p_s(c_s, \boldsymbol{p}_s) + \big(\sigma + (1-\sigma)\, \xi_r\big)\, \log c_s. \tag{78}$$

The partial derivative with respect to $\log\, p_{krs}$ is then given as

$$\frac{\partial \log c_{krs}(c_s, \boldsymbol{p}_s)}{\partial \log p_{krs}} = -\sigma + \sigma \frac{\partial \log p_s(c_s, \boldsymbol{p}_s)}{\partial \log p_{krs}} + \Big(\sigma + (1-\sigma)\, \xi_q\Big) \frac{\partial \log c_s}{\partial \log p_{krs}}. \tag{79}$$

In order to see how the nonhomothetic ideal price-index p_s responds to a *ceteris paribus* change in p_{krs}, I compute

$$\begin{aligned} \frac{\partial p_s(c_s, \boldsymbol{p}_s)}{\partial p_{krs}} &= \frac{\partial}{\partial p_{krs}} \left[\left(\sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} \varphi_q\, p_{iqs}^{1-\sigma}\, c_s^{(1-\sigma)(\xi_q - 1)} \right)^{\frac{1}{1-\sigma}} \right] \\ &= p_s(c_s, \boldsymbol{p}_s)^{\sigma}\, \varphi_q\, p_{iqs}^{-\sigma}\, c_s^{(1-\sigma)(\xi_q-1)-1} \\ &\quad + p_s(c_s, \boldsymbol{p}_s)^{\sigma} \frac{\partial c_s}{\partial p_{krs}} \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} \varphi_q\, p_{iqs}^{1-\sigma}\, c_s^{(1-\sigma)(\xi_q-1)-1} (\xi_q - 1) \end{aligned} \tag{80}$$

Using (63) and writing the above derivative as an elasticity, we obtain

$$\frac{\partial\, \log p_s(c_s, \boldsymbol{p}_s)}{\partial\, \log p_{kqs}} = x_{kqs}(c_s, \boldsymbol{p}_s) + \frac{\partial \log c_s}{\partial \log p_{kqs}} \Big[\bar{\xi}_s(c_s, \boldsymbol{p}_s) - 1 \Big] \tag{81}$$

Next, $\partial \log c_s / \partial \log p_{kqs}$ is most conveniently recovered in an application of the

implicit function theorem. Specifically, I define

$$F(c_s, p_{kqs}) \equiv \sum_{q=1}^{Q} \sum_{i=1}^{N_{qs}} p_{iqs} c_{iqs}(c_s, \boldsymbol{p}_s)\, \xi_q - A \left(\frac{c_s}{c}\right)^{\frac{\eta-1}{\eta}} \tag{82}$$

We can, therefore compute

$$\begin{aligned} \frac{\partial F(c_s, p_{kqs})}{\partial c_s} = \; & \sigma \sum_{q=1}^{Q} \varphi_q\, \xi_q \sum_{i=1}^{N_{qs}} \left(\frac{p_{iqs}}{p_s(c_s, \boldsymbol{p}_s)}\right)^{1-\sigma} \frac{\partial p_s(c_s, \boldsymbol{p}_s)}{\partial c_s}\, c_s^{\sigma+(1-\sigma)\,\xi_q} \\ & + \sum_{q=1}^{Q} \varphi_q\, \xi_q \left(\sigma + (1-\sigma)\,\xi_q\right) \sum_{i=1}^{N_{qs}} \left(\frac{p_{iqs}}{p_s(c_s, \boldsymbol{p}_s)}\right)^{1-\sigma} p_s(c_s, \boldsymbol{p}_s)\, c_s^{(1-\sigma)(\xi_q-1)} \\ & - A\, \frac{\eta-1}{\eta} \left(\frac{c_s}{c}\right)^{\frac{\eta-1}{\eta}-1} \end{aligned} \tag{83}$$

as well as

$$\begin{aligned} \frac{\partial F(c_s, p_{kqs})}{\partial p_{krs}} = & \varphi_r\, \xi_r\, (1-\sigma) \left(\frac{p_{krs}}{p_s(c_s, \boldsymbol{p}_s)}\right)^{-\sigma} c_s^{\sigma+(1-\sigma)\,\xi_r} \\ & + \sigma \sum_{q=1}^{Q} \varphi_q\, \xi_q \sum_{i=1}^{N_{qs}} \left(\frac{p_{iqs}}{p_s(c_s, \boldsymbol{p}_s)}\right)^{1-\sigma} \left.\frac{\partial p_s(c_s, \boldsymbol{p}_s)}{\partial p_{krs}}\right|_{c_s} c_s^{\sigma+(1-\sigma)\,\xi_q} \end{aligned} \tag{84}$$

Note that, in the partial derivative of F with respect to p_{krs} sectoral consumption c_s is not responsive to price changes. That is,

$$\left.\frac{\partial p_s(c_s, \boldsymbol{p}_s)}{\partial p_{krs}}\right|_{c_s} \neq \frac{\partial p_s(c_s, \boldsymbol{p}_s)}{\partial p_{krs}} \tag{85}$$

and, specifically,

$$\left.\frac{\partial p_s(c_s, \boldsymbol{p}_s)}{\partial p_{krs}}\right|_{c_s} = x_{krs}(c_s, \boldsymbol{p}_s). \tag{86}$$

By the implicit function theorem (and a few straightforward algebraic manipula-

tions), we then have

$$\frac{\partial \log c_s}{\partial \log p_{krs}} = -\frac{\dfrac{\partial F(c_s, p_{krs})}{\partial p_{krs}} \dfrac{p_{krs}}{p_s(c_s, \boldsymbol{p}_s)}}{\dfrac{\partial F(c_s, p_{krs})}{\partial p_{krs}} \dfrac{c_s}{p_s(c_s, \boldsymbol{p}_s)}} \tag{87}$$

Substituting in the results from above and simplifying the expression we find that

$$\frac{\partial \log c_s}{\partial \log p_{krs}} = -\frac{(1-\sigma)\, x_{krs}(c_s, \boldsymbol{p}_s)\, \xi_r + \sigma\, x_{krs}(c_s, \boldsymbol{p}_s)\, \bar{\xi}_s(c_s, \boldsymbol{p}_s)}{(1-\sigma)\, \bar{\xi}_s^2(c_s, \boldsymbol{p}_s) + \sigma\, \bar{\xi}_s(c_s, \boldsymbol{p}_s)^2 - \dfrac{\eta - 1}{\eta}\, \bar{\xi}_s(c_s, \boldsymbol{p}_s)}. \tag{88}$$

Finally, substituting (81) and (88) into (79) and by duality, the price-elasticity of variety (k, r, s) is given as

$$\left| \frac{\partial \log c_{krs}(y, \boldsymbol{p})}{\partial \log p_{krs}} \right| = \big(1 - x_{krs}(y, \boldsymbol{p})\big)\, \sigma + x_{krs}(y, \boldsymbol{p})\, \eta\, \zeta_{krs}(y, \boldsymbol{p}) \tag{89}$$

where

$$\zeta_{krs}(y, \boldsymbol{p}) \equiv \frac{\Big(\sigma\, \bar{\xi}(y, \boldsymbol{p}) + (1-\sigma)\, \xi_r\Big)^2}{\sigma\, \eta\, \bar{\xi}(y, \boldsymbol{p})^2 + (1-\sigma)\, \eta\, \bar{\xi^2}(y, \boldsymbol{p}) + (1-\eta)\, \bar{\xi}(y, \boldsymbol{p})}. \tag{90}$$

Algorithm to compute Nash equilibrium.

1. For each $y \in \text{support}(g)$ guess $\{y_s^{(0)}\}$ such that $\int y_s^{(0)}\, ds = y$

2. In iteration n, conditional on $\{y_s^{(n)}\}$, find $\{\boldsymbol{\mu}_s^{(n)}\}$ such that equations (4) hold for

$$\boldsymbol{p}_s^{(n)} = \boldsymbol{\mu}_s^{(n)} \circ \boldsymbol{\lambda}_s$$

   This is an isolated $Q \times N$ dimensional root-finding problem for each $s \in \mathscr{S}$.

3. Compute $c_{iqs}^{(n)} = c_{iqs}(y_s^{(n)}, \boldsymbol{\mu}_s^{(n)} \circ \boldsymbol{\lambda}_s)$.

4. Find $\{y_s^{(n+1)}\}$ and $c^{(n+1)}$ such that

$$\frac{\sum_q \sum_i \mu_{iqs}^{(n)} \lambda_{iqs} c_{iqs}^{(n)} \xi_q}{\int \sum_q \sum_i \mu_{iqs}^{(n)} \lambda_{iqs} c_{iqs}^{(n)} \xi_q \, ds} = \left(\frac{y_s^{(n+1)}}{p_s^{(n)} c^{(n+1)}} \right)^{\frac{\eta-1}{\eta}} \quad \forall \quad s \in \mathscr{S}$$

and

$$\int y_s^{(n+1)} \, ds = y$$

5. Return to 2. and iterate until $\left|\left| \boldsymbol{\mu}_s^{(n+1)} - \boldsymbol{\mu}_s^{(n)} \right|\right| <$ tolerance.