

Inequality*

Gideon Bornstein[†]

Alessandra Peter[‡]

March 26, 2026

Abstract

This paper studies the distributional effects of market power when firms endogenously choose their product portfolios. We develop a general equilibrium model in which oligopolistic firms choose the quality of their products and compete both with other firms and within their own product lines. High-market-power firms distort their low-end products to limit cannibalization of their high-end sales, whereas low-market-power firms face too much external competition for such distortions to arise. We show that an increase in market power disproportionately harms low-income households by lowering the quality of the goods they consume and raising their unit prices. By contrast, through general equilibrium forces, high-income households may benefit from weaker competition. We test a key prediction of the model using data on grocery store purchases and find that in more concentrated markets, dominant firms' market shares tilt toward high-income consumers. In a quantitative version of the model calibrated to the U.S. economy, market power reduces the welfare of low-income households by 4.8% relative to the efficient allocation and increases that of high-income households by 2.0%. The rise in concentration over the past four decades widened the welfare gap between rich and poor by 0.9 percentage points.

*We thank Pete Klenow, Virgiliu Midrigan and Ezra Oberfield for valuable comments and suggestions, as well as seminar participants at the SED, Normac, AEA, Wharton, NYU, Boston University, Boston College, Notre Dame, and Cornell for helpful feedback. Jacky Chen provided excellent research assistance. 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.

[†]The Wharton School, University of Pennsylvania and NBER; gideonbo@wharton.upenn.edu

[‡]Department of Economics, New York University, NBER, and CEPR; alessandra.peter@nyu.edu

1 Introduction

Over the past four decades, U.S. product markets have grown substantially more concentrated. The employment share of the largest firms has increased across most sectors (Autor et al., 2020) and so have markups (De Loecker et al., 2020). A large literature has studied the aggregate efficiency consequences of this rise in market power. But market power may also have *distributional* consequences. Rich and poor households often purchase different products from the same firm: a Toyota Camry or a Lexus, a coach seat or business class, an iPhone SE or an iPhone Pro Max. How firms design these product portfolios, and the prices at which they offer them, depends on the competitive pressure they face from rival firms. This paper studies how market power shapes the product design choices of firms and, through this channel, affects consumption inequality.

To study the distributional consequences of market power, we proceed in three steps. First, we develop a general equilibrium model in which oligopolistic firms endogenously choose the quality and prices of their products. Firms compete both externally with other firms in the sectors and internally through cannibalization across their own product lines. We show that only high-market-power firms have cannibalization concerns: these firms distort downward the quality of their low-end goods, raise their low-end prices, and lower the prices of their high-end products. By contrast, low-market-power firms, which face sufficient external competition, have no incentive to distort their product design. As a result, an increase in market power hurts low-income households disproportionately, while general equilibrium forces can even make high-income households better off in absolute terms. Second, we provide empirical evidence for the model’s key prediction using data on grocery store goods, showing that in more concentrated product markets, the largest firms sell disproportionately more to high-income consumers. Third, we calibrate the model to quantify the distributional effects of market power. We find that, relative to the efficient allocation, market power lowers the welfare of low-income households by 4.8% while raising that of high-income households by 2.0%.

The key mechanism behind the distributional impact of market power is *cannibalization-driven quality distortion*. When a firm with market power sells products to consumers of different income levels, it faces an internal trade-off: making the low-end product too attractive risks cannibalizing sales of its high-end offering. To prevent high-income consumers from downgrading, the firm deliberately reduces the quality of its low-end product below the efficient level and raises its price. This mechanism is muted with sufficient competition, as firms that degrade their products would simply lose customers to competitors. Our model extends Mussa and Rosen (1978), who study endogenous quality choices of a monopolist, to an environment with oligopolistic competition in general equilibrium.

The model we construct features a continuum of sectors, each populated by a finite number of firms that differ in their brand value. Consumers differ in income and, within each sector, choose a single product from a single firm. Each firm designs a menu of vertically differentiated products from which consumers self-select according to their income. There are two sources of market power in our model: the number of firms in a sector and the dispersion in brand values across firms. The model’s structure allows us to isolate each source and to decompose the welfare costs of market power into those that fall symmetrically on all households and those that disproportionately burden the poor.

We develop our argument in two steps. We start by analyzing an “island economy” in which high- and low-income consumers are segregated into separate markets, so that cannibalization concerns are absent. In this benchmark, imperfect competition still generates misallocation across firms: high-brand-value firms charge higher markups and attract too few consumers, a distortion in the spirit of [Atkeson and Burstein \(2008\)](#). However, product design remains efficient. All firms offer quality levels that equate the marginal utility of quality to its marginal cost of production, just as the social planner would. As a result, consumption inequality equals income inequality, regardless of the degree of market power. The island economy isolates the standard welfare cost of oligopolistic competition: prices that are too high and too little consumption from the largest firms. This welfare cost, while present, falls equally on rich and poor households.

In the full market equilibrium, high- and low-income consumers are no longer segregated, and the island allocation may not be sustainable. High-brand-value firms charge higher markups in the island economy, so the price gap between their high- and low-end products is large. If the markup is large enough, high-income consumers at these firms would prefer to purchase the cheaper low-end product and use the savings in other sectors. To prevent this, firms with sufficient market power distort the quality of their low-end product downward and raise its price, while reducing the price of their high-end product. In general equilibrium, the resources freed by the reduction in low-end quality must be absorbed elsewhere. Some of these resources are redirected toward the goods consumed by high-income households, whose equilibrium quality rises above the efficient level. Through this general equilibrium channel, high-income households benefit from the very distortions that hurt the poor.

There are two sources of market power in the model: the number of competing firms and the dispersion in brand value across firms. We show that both sources increase welfare inequality. When the number of firms falls, each firm’s market share rises, leading to an increase in desired markups. With higher markups, there is more scope for cannibalization and firms distort quality at the bottom more. When brand-value dispersion increases, “superstar” firms emerge with disproportionately large market shares. These firms face the strongest cannibalization incentives, so they distort their low-end products the most and reduce prices of high-end products. As a result, misallocation of consumers across firms worsen for low-income consumers who are pushed toward lower-brand-value firms and improves for high-income ones, who gravitate toward the now-cheaper high-end products.

The model delivers a sharp empirical prediction: in more concentrated markets, dominant firms should have a higher market share among high-income households relative to low-income ones. We test this prediction using the NielsenIQ Consumer Panel, which records household-level purchases of grocery products along with household demographics. For each product market, we construct the relative market share of the top firm across income groups. We find that in more concentrated product markets, the top firm’s relative market share tilts more toward the high-income segment. In the typical market, the top firm’s market share among high-income households is about 4 log points higher than among low-income ones. A one standard deviation increase in concentration, as measured by the Herfindahl-Hirschman Index (HHI), more than doubles this gap. This relationship holds when we use alternative measures of concentration, or when we broaden the definition of top firms to the largest three.

To quantify the distributional effects of market power, we calibrate the model to the 2022 U.S. economy, targeting spending inequality from the PSID, the average number of firms per 6-digit NAICS industry, and the average HHI across these industries. In the calibrated model, market power reduces the welfare of low-income households by 4.8% relative to the efficient allocation, while high-income households are 2.0% better off. To isolate the role of product design, we compare these results to the island economy, which features the same degree of imperfect competition but shuts down cannibalization across income groups. There, the welfare loss is a symmetric 0.3% for both income groups. The distributional costs of market power are an order of magnitude larger than the standard efficiency costs. They are also large enough to flip the sign for high-income households: the gains from lower high-end prices and higher quality more than offset the standard welfare costs of imperfect competition.

We then use the model to assess the distributional impact of the rise in market concentration since 1980. We model this rise as an increase in the dispersion of brand values across firms, consistent with the emergence of dominant superstar firms. The increase in concentration over this period widened the welfare gap between high- and low-income households by 0.9 percentage points. Our analysis shows that less competition, even if driven by the emergence of superstar firms and not merely through a reduction in the number of competitors, leads to more product distortion and, through it, higher consumption inequality.

Related literature. This paper connects to several strands of work.

First, it contributes to the literature on the aggregate consequences of rising market power. A large literature has studied the efficiency costs of markups while abstracting from distributional consequences, e.g., (De Loecker et al., 2020; Edmond et al., 2022; Baqaee and Farhi, 2020; Hsieh and Klenow, 2009; Autor et al., 2020). A narrower literature has focused on the distributional consequences of market power through its effect on factor income shares (Karabarbounis and Neiman, 2014; Boar and Midrigan, 2024). Our paper shows that market power also has distributional consequences that operate through a distinct channel: the endogenous design of product portfolios.

Second, our paper is related to the theoretical literature on product quality distortions. Starting with Mussa and Rosen (1978), this literature shows that a monopolist may distort the quality offered to low-valuation consumers. Unlike that canonical framework, firms in our setting compete for customers with other firms in the sector. Two related papers that consider oligopolistic competition are Stole (1995) and Rochet and Stole (2002), both of which show that intensive competition reduces the scope of product quality distortions. Our analysis adds to these papers by developing a tractable general equilibrium framework which allows us to study how market structure shapes the distribution of consumption across income groups and to quantify aggregate welfare effects. We show that, unlike in the partial-equilibrium comparative statics, general equilibrium forces can imply that high-income households benefit from a lack of competition.

Finally, the paper is related to a literature studying how prices, markups, and product availability vary with household income. Jaravel (2019) and Oberfield (2023) study whether product innovations disproportionately benefit high- or low-income consumers. Kaplan and Menzio (2015), Nord (2023), and Sangani (2024) document large price dispersion across stores and relate it to household search

behavior. [Becker \(2024\)](#) and [Mongey and Waugh \(2025\)](#) study the role of non-homothetic preferences in firm pricing. Relative to this literature, we endogenize the set of products that are offered to consumers and find that distortions in quality hurt low-income consumers much more than markups.

Layout. The paper proceeds as follows. Section 2 presents the model, defines the competitive equilibrium, and characterizes the efficient allocation. Section 3 analyzes the island economy and the market equilibrium, and derives comparative statics linking market power to consumption inequality. Section 4 presents empirical evidence. Section 5 describes the calibration and quantitative results. Section 6 concludes.

2 Model

This section develops a general-equilibrium model in which firms endogenously choose the quality levels of the goods they sell. Households differ in their income levels and purchase a single good in each sector. We first describe preferences and household demand, then turn to firm behavior, and finally define the competitive equilibrium.

2.1 Environment

Households. The economy is populated by two types of households, low- and high-income ones. There is a measure one-half of each type of household. Labor supply is inelastic. High-income households have high labor endowment $L_h = 1 + \alpha$ and receive a higher share of aggregate profits $(1 + \alpha)$. Low-income households have low labor endowment $L_l = 1 - \alpha$ and receive the remainder $(1 - \alpha)$ of aggregate profits. We denote households by $i \in (0, 1)$, where households in $(0, \frac{1}{2})$ are low-income and households in $(\frac{1}{2}, 1)$ are high-income.

Households consume products from a continuum of sectors, indexed by $s \in (0, 1)$. The aggregate consumption is a Cobb-Douglas aggregator across sectoral consumption,

$$\ln C_i = \int_0^1 \ln c_i(s) ds, \tag{2.1}$$

where $c_i(s)$ is the utility-adjusted consumption of household i from sector s . Within each sector, there is a finite number of firms, N , each offering a menu of vertically differentiated products. Products differ in the quality level q_{jk} , where jk is the k -th product sold by firm j . Households face a discrete choice problem within each sector, i.e., they choose to purchase a single product k from a single firm j .

Utility-adjusted consumption consists of three components. First, households derive utility from a product's quality. Second, households have preferences towards the products sold by firm j . This captures both brand value and non-rival product features, such as access to the iOS operating system. We assume that these come in the form of taste shifters and have a common component η_j and an idiosyn-

cratic component $\varepsilon_{ij}/(\sigma - 1)$.¹ The idiosyncratic taste shifter ε_{ij} is independent across households and sectors and follows a standard Gumbel distribution. This assumption allows us to characterize the demand function of households analytically. Formally, the utility-adjusted consumption from product jk is given by

$$\ln c_{ijk} = \ln q_{jk} + \eta_j + \frac{\varepsilon_{ij}}{\sigma - 1}. \quad (2.2)$$

Households solve a collection of discrete choice problems: In each sector s , they choose the firm $j(s)$ and product $k(s)$ to maximize their utility, subject to their type's budget constraint.

$$\begin{aligned} \max_{\{j(s), k(s), c_i(s)\}} & \int \ln c_i(s) ds \\ \text{s.t.} & \int p_{j(s)k(s)} ds = E_i, \\ & \ln c_i(s) = \ln q_{j(s)k(s)} + \eta_{j(s)} + \frac{\varepsilon_{ij(s)}}{\sigma - 1} \end{aligned} \quad (2.3)$$

Let $1/P_i$ denote the inverse Lagrange multiplier of the problem above. That is, P_i measures the cost of acquiring an additional unit of utility for household i . We can use the marginal price index P_i to characterize the optimal discrete choice in each sector. Below, we omit the sector subscript for ease of notation. Household i chooses the product k from firm j that maximizes

$$\{j_i, k_i\} = \operatorname{argmax}_{\{j, k\}} \ln(q_{jk}) - \frac{p_{jk}}{P_i} + \eta_j + \frac{\varepsilon_{ij}}{\sigma - 1}. \quad (2.4)$$

As we prove in the Appendix, the continuum of sectors implies that idiosyncratic tastes do not result in ex-post heterogeneity in overall consumption conditional on household income. In particular, we denote by P_l and P_h the marginal price index of low- and high-income households, respectively.

Firms. Sectors are symmetric. In each sector s , there is a fixed number N of firms. Firms differ in their brand value η_j , but are otherwise identical. The cost of producing a good is linear in the quality offered: producing a good of quality q costs κq units of labor. Each firm decides the qualities and prices of the products it sells. We restrict attention to deterministic mechanisms and assume that firms must offer their menu of products to all customers.² Thus, firms must design their product menu such that each consumer self-selects into their designated bundle.

There are two dimensions of consumer heterogeneity. First, consumers differ in their income level, which affects their marginal price index. Second, consumers differ in their idiosyncratic taste for brands. Idiosyncratic taste shocks affect which firm households choose to buy from, but do not affect their valuation of quality conditional on purchasing from a given firm. As a result, firms optimally offer only two products $k \in \{l, h\}$: one for high-income, and one for low-income consumers.

¹The parameter σ , as we show below, governs the degree of substitutability across firms.

²The restriction to deterministic mechanisms implies firms are not allowed to sell lottery tickets to obtain their products.

Let q_l be the product the firm offers to low-income consumers and p_l its price. $\{q_h, p_h\}$ denotes the corresponding product and price for high-income consumers. The profit maximization problem of a firm with brand value η_j is given by

$$\begin{aligned} \pi_j = & \max_{\{q_l, q_h, p_l, p_h, B_{jl}, B_{jh}\}} B_{jl} (p_l - \kappa q_l) + B_{jh} (p_h - \kappa q_h), \\ \text{s.t.} & B_{jl} = \mathcal{B}(q_l, p_l, \eta_j; \{q_{j'l}, p_{j'l}, \eta_{j'}\}_{j' \neq j}, P_l), \\ & B_{jh} = \mathcal{B}(q_h, p_h, \eta_j; \{q_{j'h}, p_{j'h}, \eta_{j'}\}_{j' \neq j}, P_h), \\ & \ln(q_h) - \ln(q_l) \geq \frac{p_h - p_l}{P_h} \\ & \ln(q_h) - \ln(q_l) \leq \frac{p_h - p_l}{P_l} \end{aligned} \quad (2.5)$$

where the first two constraints are the customer bases as a function of product offering and market conditions, and the last two are the incentive compatibility of high- and low-income households, respectively.

The incentive compatibility constraints ensure each customer is better off choosing their bundle. The high-income consumer must prefer purchasing the high-quality product to purchasing the low-quality one and using the extra funds in other sectors. The value of additional funds is given by the inverse of the marginal price index P_h . Similarly, the low-income household does not find it beneficial to switch to the high-quality product and pay the price difference, which they value at P_l .

Leveraging the Gumbel distributional assumption, the following Lemma derives an analytic characterization for the customer base functions.

LEMMA 1. *The customer base function is given by*

$$\mathcal{B}(q_j, p_j, \eta_j; \{q_{j'}, p_{j'}, \eta_{j'}\}_{j' \neq j}, P) = \frac{e^{(\sigma-1)\left(\ln(q_j) - \frac{p_j}{P} + \eta_j\right)}}{\sum_{j'} e^{(\sigma-1)\left(\ln(q_{j'}) - \frac{p_{j'}}{P} + \eta_{j'}\right)}}. \quad (2.6)$$

2.2 Equilibrium

We restrict attention to a symmetric equilibrium. An equilibrium is a set of firm-level prices and product qualities $\{p_{jl}, p_{jh}, q_{jl}, q_{jh}\}_j$, market share distributions $\{B_{jl}, B_{jh}\}_j$, firm-level and aggregate profits $\{\{\pi_j\}_j, \Pi\}$, income levels $\{E_l, E_h\}$, and marginal price indices $\{P_l, P_h\}$, such that

1. Given their income E_i , realization of taste shocks ε_{ij} , and prices and qualities offered, each household chooses the optimal firm and quality level in each sector, as defined by (2.4), and the household problem (2.3) is solved.
2. Given the marginal price indices P_h and P_l as well as other firms product offerings and prices $\{p_{j'l}, p_{j'h}, q_{j'l}, q_{j'h}\}_{j' \neq j}$, product offerings and prices solve the firms' problem (2.5).
3. The market shares $\{B_{jl}, B_{jh}\}_j$ are consistent with the market share function (2.6).

4. The marginal price indices are defined as the inverse Lagrange multipliers on the budget constraints in the household problem (2.3).
5. Aggregate profits in the economy are given by

$$\Pi = \sum_j \pi_j. \quad (2.7)$$

6. The income of each household is equal to their labor earnings plus their share of aggregate profits:

$$\begin{aligned} E_l &= (1 - \alpha)(1 + \Pi), \\ E_h &= (1 + \alpha)(1 + \Pi). \end{aligned}$$

7. The labor market clears:

$$\kappa \sum_j [B_{jl}q_{jl} + B_{jh}q_{jh}] = 1. \quad (2.8)$$

2.3 Welfare, Discrete Choice, and the Entropy Measure of Concentration

We now discuss the determinants of welfare in the equilibrium allocation. Because there is a continuum of symmetric sectors, the idiosyncratic taste shocks wash out in the aggregate, and all households of the same income level have the same level of aggregate consumption. Let $\tau \in \{l, h\}$ denote the income-type of households. The following proposition derives aggregate consumption of a household as a function of the quality, brand value, and market share of each firm.

PROPOSITION 1. *Suppose households optimally choose which brand to purchase, according to (2.4). Then, the welfare of a household of type τ is given by*

$$\ln(C_\tau) = \sum_j B_{j\tau} (\ln q_{j\tau} + \eta_j) + \frac{\sum_j B_{j\tau} \ln(1/B_{j\tau})}{\sigma - 1} + \frac{\gamma}{\sigma - 1} \quad (2.9)$$

where $B_{j\tau}$ is the fraction of markets in which household τ purchases from firm j , and γ is the Euler-Mascheroni constant.

Proposition 1 shows that aggregate consumption is the sum of three terms. The first term in the RHS of equation (2.9) is a weighted average of the quality and brand value of the products a household consumes, weighted by $B_{j\tau}$, the share of markets in which firm j is chosen.

The second and third terms in equation (2.9) capture the utility from the idiosyncratic taste shocks. The third term equals the utility a household would derive from the idiosyncratic taste shocks if there were only a single firm— γ is the mean of the standard Gumbel distribution. Our preferences feature love-of-variety: more choices allow the household to choose a product towards which they have a higher taste. How much the availability of an additional brand increases consumer welfare depends on how often it is actually chosen, which is summarized by its market share.

Proposition 1 shows that a sufficient statistic capturing the utility derived from idiosyncratic taste shocks is the entropy measure of concentration, $\sum_j B_{j\tau} \ln(1/B_{j\tau})$. This concentration measure was introduced in Finkelstein and Friedberg (1967), and to our knowledge, ours is the first paper that derives a micro-foundation for its direct effect on welfare. The term “entropy measure of concentration” may be misleading, as higher concentration reduces this measure. Rather, it is more appropriate to call it the “entropy measure of competition.” When there is only one firm to choose from, the entropy measure equals zero. A larger number of firms or lower dispersion in brand value increases the entropy measure and increases the expected idiosyncratic taste shock from the chosen brand.

2.4 Efficient allocation

We next turn to derive the efficient allocation by solving the problem of a social planner who chooses allocations subject to the same production technology and the same information structure. To abstract from distributional considerations, we impose an additional constraint on the social planner: all consumption by low-income households has to be produced using labor of low-income households and vice versa.³

$$\begin{aligned}
& \max_{\{j_{i\tau}(s), q_{i\tau}(s)\}_{\{i,\tau,s\}}} \int_i \int_s \ln(c_i(s)) ds di & (2.10) \\
& \text{s.t.} \quad \ln(c_i(s)) = \ln(q_i(s)) + \eta_{j_i(s)} + \frac{\varepsilon_{ij_i(s)}}{\sigma - 1}, \quad \forall i, s \\
& \int_0^{\frac{1}{2}} \int_s \kappa q_i(s) ds di = L_l, \\
& \int_{\frac{1}{2}}^1 \int_s \kappa q_i(s) ds di = L_h, \\
& \int_s \ln(c_i(s)) ds \geq \int_s \ln(c_{i'}(s)) ds, \quad \forall i, i' \in (0, \frac{1}{2}), \\
& \int_s \ln(c_i(s)) ds \geq \int_s \ln(c_{i'}(s)) ds, \quad \forall i, i' \in (\frac{1}{2}, 1),
\end{aligned}$$

where the first constraint is the utility-adjusted consumption, the second set of constraints is the resource for high- and low-income types respectively, and the third set of constraints is the incentive compatibility constraint, which ensures each household happily consumes their designated bundle. Note that the incentive compatibility constraint only considers deviating to alternative bundles *within* income type, as the planner can infer the income type of consumers.⁴

Let P_h^{FB} and P_l^{FB} denote the inverse Lagrange multipliers on the resource constraints for the high and low types respectively.

PROPOSITION 2. *In the efficient allocation, the incentive compatibility constraints are slack, and the allocation features*

³This assumption is analogous to choosing a specific set of Pareto-weights in the problem of a utilitarian social planner.

⁴This could easily be implemented by offering a bundle that contains not only the consumption basket but also the efficient units of labor provided by the household.

(i) quality levels that are identical across firms and are given by

$$q_{\tau}^{FB} = \frac{L_{\tau}}{\kappa}, \quad \forall \tau \in \{l, h\}, \quad (2.11)$$

(ii) market shares of each firm j that are identical across the two types of consumer segments and are given by

$$B_{j\tau}^{FB} = \frac{e^{(\sigma-1)\eta_j}}{\sum_{j'} e^{(\sigma-1)\eta_{j'}}}, \quad \forall \tau \in \{l, h\}. \quad (2.12)$$

All proofs are relegated to Appendix A.

Proposition 2 shows that only two quality levels are offered by the planner, one for all high-income and one for all low-income consumers. This is a result of two features of the model. First, all firms have the same cost of producing quality, κ . Second, the only household heterogeneity that affects the marginal valuation of quality is income, and hence there is no incentive to tailor quality levels to a household's realization of taste shocks ε_{ij} .

Each of the two quality levels equates marginal utility to the marginal cost of production. Since low-income households have a lower labor endowment, their real marginal cost κ/L_l is higher than the real marginal cost of high-income households, κ/L_h and the social planner optimally produces lower-quality goods for low-income people.

Since production costs and quality levels are the same for all firms, the planner allocates households to firms according to their brand value, taking into account both the common component η_j as well as each household's idiosyncratic tastes ε_{ij} . Since the planner allocates based on household tastes, households have no incentive to misreport their idiosyncratic taste shifters, and incentive compatibility constraints are slack. Each firm's market share is simply a function of its brand value η_j relative to its competitors'—idiosyncratic taste shocks are i.i.d. and wash out across the continuum of consumers.

In the efficient allocation, the market share of firm j is identical across the low- and high-income segments of the population. The only difference across income types in the efficient allocation is the quality levels offered, which depend on each household's labor income. Therefore, consumption inequality is simply equal to income inequality:

LEMMA 2. *In the efficient allocation, consumption inequality is equal to income inequality. That is,*

$$\ln C_h - \ln C_l = \ln L_h - \ln L_l = \ln \left(\frac{1 + \alpha}{1 - \alpha} \right). \quad (2.13)$$

3 Distortions in product design, misallocation, and market power

In this section, we analyze the market equilibrium and discuss sources of inefficiencies in the model.

We do so in two steps. We start by considering an *island economy* in which high- and low-income consumers are segregated and hence consumer heterogeneity has no direct impact on firm choices.

In this economy, we show that there is misallocation across firms due to imperfect competition: the market share of high-quality firms is too low. However, product design remains efficient and there are no differential distortions across income groups: consumption inequality is identical to the efficient allocation.

We then characterize the full market equilibrium. In the absence of sufficient competition, firms *distort* their product design choices to ensure that consumers self-select into their respective allocation. Low-income consumers are sold products of too low quality, while high-income consumers are sold products of too high quality. Further, low-income consumers are disproportionately excluded from high-quality firms. Taken together, consumption inequality is higher than in the efficient allocation.

3.1 Island economy

We start by analyzing the island economy. The firm's problem is identical to (2.5) but without the last two constraints that ensure incentive compatibility. Hence, firms effectively solve two separate profit-maximization problems, one for each consumer type. The problem of firm j in the island of consumers of type τ is given by

$$\begin{aligned} \max_{\{p_{j\tau}, q_{j\tau}, B_{j\tau}\}} & B_{j\tau}(p_{j\tau} - \kappa q_{j\tau}), \\ \text{s.t.} & B_{j\tau} = \frac{e^{(\sigma-1)(\ln(q_{j\tau}) - \frac{p_{j\tau}}{P_\tau} + \eta_j)}}{\sum_{j'} e^{(\sigma-1)(\ln(q_{j'\tau}) - \frac{p_{j'\tau}}{P_\tau} + \eta_{j'})}}. \end{aligned} \quad (3.1)$$

Optimal product quality $q_{j\tau}$ and price $p_{j\tau}$ are given by

$$q_{j\tau} = \frac{P_\tau}{\kappa} \quad (3.2)$$

$$p_{j\tau} = \kappa q_{j\tau} \left(\frac{\sigma}{\sigma-1} + \frac{1}{\sigma-1} \frac{B_{j\tau}}{1-B_{j\tau}} \right). \quad (3.3)$$

Firms choose the quality level to maximize consumer surplus, equating the marginal utility ($\frac{1}{q_{j\tau}}$) to the real marginal cost of producing quality ($\frac{\kappa}{P_\tau}$). Since taste shocks are additive, this level of quality is the same for households of the same income level.

When choosing its price, the firm trades off attracting additional customers through a low price against profits per customer. The optimal markup set by the firm, $\frac{p_{j\tau}}{\kappa q_{j\tau}}$, consists of two components. The first component is the classic monopolistic markup $\frac{\sigma}{\sigma-1}$ that arises when firms are atomistic. The second component depends positively on the market share of the firm in that consumer segment, $B_{j\tau}$. As a firm's market share grows, the potential customers it can attract by lowering its price, $1 - B_{j\tau}$, shrinks. As a result, firms with larger market shares face smaller demand elasticities, and they charge a higher markup.

To close the economy, we maintain the assumption of a single labor market (Equation 2.8) and that aggregate profits (across the two islands) are split between high- and low-income consumers in

proportion to their labor endowment. The following proposition characterizes the resulting equilibrium allocation in the island economy.

PROPOSITION 3. *There exists a unique equilibrium in the island economy whose allocation features*

(i) *quality levels that are identical across firms and are given by*

$$q_{\tau}^{IE} = \frac{L_{\tau}}{\kappa}, \quad \forall \tau \in \{l, h\}, \quad (3.4)$$

(ii) *market shares of each firm j that are identical across the two types of consumer segments and are defined implicitly by*

$$B_{j\tau}^{IE} = \frac{e^{(\sigma-1)\eta_j - \frac{1}{1-B_{j\tau}^{IE}}}}{\sum_{j'} e^{(\sigma-1)\eta_{j'} - \frac{1}{1-B_{j'\tau}^{IE}}}}, \quad \forall \tau \in \{l, h\}. \quad (3.5)$$

Proposition 3 shows that quality levels in the island economy are identical to the efficient allocation characterized in Proposition 2. The firms' optimality condition (3.2) implies that all firms sell the same level of quality on an island. Proposition 3 further establishes that there is no cross-subsidization in the island economy, so that the quality levels are not only constant within each island but also identical to the efficient allocation.

The second part of Proposition 3 shows that market shares of firms are, in general, different from the efficient allocation. Since high-brand-value firms charge more for the same quality level, their market share is inefficiently low. The following corollary formalizes this result.

COROLLARY 1. *There is misallocation of consumers across firms in the island economy, with high-brand-value firms serving too few, and low-brand-value firms serving too many consumers. Formally, for any two firms j and k , with $\eta_j > \eta_k$, we have*

$$\frac{B_{j\tau}^{IE}}{B_{j\tau}^{FB}} < \frac{B_{k\tau}^{IE}}{B_{k\tau}^{FB}}. \quad (3.6)$$

This result is akin to the distortion resulting from oligopolistic competition in Atkeson and Burstein (2008). Larger firms face lower demand elasticities, charge higher markups, and end up selling less than the planner would like them to. While misallocation reduces aggregate consumption for each type of household, the following Proposition shows that this proportional reduction in consumption is identical across the two consumer types. That is, the distortion in market shares is *independent* of consumer income: In both the high-and low-income islands, market shares and hence misallocation across consumers are identical.

PROPOSITION 4. *Consumption is inefficiently low in the island economy, but consumption inequality*

is identical to consumption inequality in the efficient allocation. That is,

$$\ln C_\tau^{IE} \leq \ln C_\tau^{FB}, \quad \forall \tau \in \{l, h\} \quad (3.7)$$

$$\ln C_h^{IE} - \ln C_l^{IE} = \ln C_h^{FB} - \ln C_l^{FB}. \quad (3.8)$$

Towards a market equilibrium. In the island economy, we relaxed the two incentive compatibility constraints:

$$\ln(q_{jh}) - \ln(q_{jl}) \geq \frac{p_{jh} - p_{jl}}{P_h} \quad (3.9)$$

$$\ln(q_{jh}) - \ln(q_{jl}) \leq \frac{p_{jh} - p_{jl}}{P_l} \quad (3.10)$$

These ensure that each consumer prefers choosing their designated product over that of the other income group. Before turning to the analysis of the market equilibrium, it is useful to examine which consumers have an incentive to purchase from the other island, and which brands they are incentivized to purchase. Note that if no incentive compatibility constraint is violated, the market equilibrium allocation is identical to the island economy one.

The following proposition characterizes which incentive compatibility constraints are violated in the island economy allocation.

LEMMA 3. *In the island economy allocation, the incentive compatibility (IC) constraints of low-income households are slack. The IC constraints of high-income households are violated for the top \tilde{N} brands. That is, if the high-income IC constraint is violated for a brand with value $\bar{\eta}$, it is violated for all brands with $\eta \geq \bar{\eta}$.*

In the island economy, firms charge positive markups. At the unit price they face, all consumers would prefer to purchase only a fraction of the quality they are offered. However, since goods are indivisible, the only possible deviation available to consumers is to buy the product designed for the other income group. Low-income consumers never want to deviate, because they would like to reduce quality and the alternative product instead offers higher quality.

High-income consumers trade off a lower cost associated with the low-quality bundle with the utility loss due to the lower quality of the good. Because quality is equated across firms within each consumer type, the utility loss associated with downgrading a product is the same across all brands. However, heterogeneous markups imply that the cost savings from downgrading a high-brand-value firm is larger. As a result, there is a cutoff brand value above which high-income households would like to downgrade their products and below which the incentive compatibility is slack.

3.2 Market allocation

We next analyze the market allocation, where firms must design their product offerings so that each consumer purchases their designated bundle. We start by characterizing firms' optimal allocations, taking price indices as given. We then turn to study the role of market power in shaping consumption inequality in two specific cases.

As Lemma 3 shows, the allocations in the island economy may violate the incentive compatibility (IC) constraints of high-income consumers for brands with sufficiently high market power. Proposition 5 characterizes the firm's optimal choices in the region in which the IC constraints are slack and the region in which the high-income IC constraint is binding.⁵

PROPOSITION 5. *The firm's optimal policy in the unconstrained region as well as the region where the IC of the high type binds is given the following expressions.*

(i) *Optimal quality levels are*

$$q_{jh} = \frac{P_h}{\kappa}, \quad (3.11)$$

$$q_{jl} = \frac{P_l}{\kappa} - \psi \left(\frac{1 - \frac{P_l}{P_h}}{\kappa B_{jl}} \right), \quad (3.12)$$

(ii) *Optimal prices are*

$$p_{jh} = \kappa q_{jh} \tilde{\mu}_{jh} - \psi \frac{1}{B_{jh}} (\tilde{\mu}_{jh} - 1), \quad (3.13)$$

$$p_{jl} = \kappa q_{jl} \tilde{\mu}_{jl} + \psi \frac{1}{B_{jl}} (\tilde{\mu}_{jl} - 1), \quad (3.14)$$

where $\psi \geq 0$ is the Lagrange multiplier on the high-income IC constraint and $\tilde{\mu}_{ij} \equiv \left(\frac{\sigma}{\sigma-1} + \frac{1}{\sigma-1} \frac{B_{j\tau}}{1-B_{j\tau}} \right)$ is the unconstrained markup.

If the IC constraint is slack ($\psi = 0$), firm j offers the same bundle as it would in the island economy allocation *conditional* on the two marginal price indices. Both high- and low-income consumers are offered a quality level that equates their marginal utility to the real marginal cost. Note that, as long as the IC binds for some firm in the economy, P_h and P_l , and hence equilibrium quality levels offered are different from the island economy allocation. When $\psi = 0$, the price a firm charges also simplifies to the familiar $p_{j\tau} = \kappa q_{j\tau} \left(\frac{\sigma}{\sigma-1} + \frac{1}{\sigma-1} \frac{B_{j\tau}}{1-B_{j\tau}} \right)$. Markups may still differ across the two consumer segments, since firms do not necessarily have equal market shares in both.

Firms for which the IC is binding cannot design their product to each income group in isolation. Instead, when designing their product portfolios, firms must take into account cannibalization considerations—how increasing the surplus for one income group affects the surplus they must provide the other income group. If the bundle offered to low-income consumers is too attractive relative to the more expensive high-quality bundle, high-income consumers will not purchase their designated product. To ensure incentive compatibility, the firm has four margins it can adjust: the two quality levels as well as prices charged.

On the quality margin, firms choose to distort only the low-quality product but not the high-quality one. The quality level of each good in the island economy was chosen to maximize the joint

⁵In theory, there could be an additional region in which the incentive compatibility constraint of low-income households is binding. However, we confirm numerically that in general equilibrium there is no parameter configuration in which that is the case for any firm.

surplus, $\ln(q_j\tau) + \eta_j - \frac{\kappa}{P_\tau}$. There is no reason for the firm to reduce the joint surplus for high-income consumers, as that would not alleviate the incentive compatibility constraint. On the other hand, reducing the surplus delivered to low-income consumers helps alleviate the IC constraint. As a result, the firm chooses to distort downwards the quality level of the low-end product.

The degree of quality distortion is governed by three variables, as shown in equation (3.12). First, the larger the Lagrange multiplier, ψ , the larger is the distortion. Second, the distortion is decreasing in the ratio of aggregate price indices, $\frac{P_l}{P_h}$. This ratio governs the relative valuation of quality of high- vs. low-income consumers. The larger is this gap in valuations, the larger the difference in willingness-to-pay for quality between the two consumer segments. So reducing quality of the low-end product hurts rich consumers more than poor ones. Third, the distortion is decreasing in the market share of low-income consumers. The higher the sales to low-income consumers, the higher the cost of distorting their quality.

While the incentive compatibility constraints push down the quality of low-quality goods produced by high-market-power firms, a general equilibrium force operates in the opposite direction. If quality levels only go down, aggregate labor demand would not meet supply, and the labor market would not clear. In general equilibrium, labor market clearing requires a decline in the real wages—the inverse of the aggregate price indices. Lower real wages, in turn, raise equilibrium quality choices uniformly across firms. As a result, the equilibrium quality of high-quality goods exceeds its first-best level. For low-quality goods, the net effect is heterogeneous: quality rises for firms with sufficiently low market power, but declines for firms with high market power.

On the price margin, firms choose to operate on both fronts: reducing the markup charged to high-income consumers and raising it for low-income consumers. Optimal prices now balance two motives. The first part of Equations (3.13) and (3.14) is the same as in the island economy: firms choose prices that trade off attracting additional customers and sales per customer. The second part of the optimal prices reflects the need to preserve incentive compatibility.

The magnitude of price distortion depends not only on the Lagrange multiplier, ψ , but also on the market share and demand elasticity in each consumer segment. The larger the market share and the demand elasticity, the larger the profit losses from price distortions. As a result, both these variables incentivize the firms to use other margins to satisfy the incentive compatibility constraint.

Relative to the island economy, welfare of each consumer type changes along two margins. First, as discussed above, high-income households get allocated a higher quality of each good they consume. Low-income households, on the other hand, are allocated, on average, less of each good. Second, firms for which the IC of the high type is binding distort prices and allocations, which results in a change in market shares. Lower prices for the rich attract more customers to these firms, while higher prices and lower quality levels for the poor lead to lower market shares in that consumer segment. Because the market share of high-market-power firms in the island economy is distorted downwards (Corollary 1), the reallocation of market shares toward high-market-power firms benefits the rich, while the reallocation of market shares away from high-market-power firms hurts the poor.

3.3 Market power and consumption inequality

There are two sources of market power in the economy: the number of competitors in a sector and the dispersion in brand value. To understand how market power shapes consumption inequality, we analyze the comparative statics to each of these sources of market power in isolation.

3.3.1 The degree of competition and the quality gap

We start by considering an economy in which there is no heterogeneity in brand value η_j across firms. In this economy, the degree of competition is solely driven by the number of firms N . Proposition 6 shows the direct link between the number of competitors and the degree of consumption inequality.

PROPOSITION 6. *When brand value is identical across firms, consumption inequality is decreasing in the number of firms N .*

With fewer firms, the market share of each, $1/N$, rises. A higher market share means firms optimally charge a higher price for their products, which increases the incentives of high-income consumers to purchase low-end products. Facing such cannibalization concerns, firms distort the quality of their low-end products downwards, increase prices at the bottom, and reduce them at the top.

Since firms are symmetric, consumers cannot substitute towards low-market power firms. Instead, their consumption decreases across the board, while the consumption of high-income households must increase in general equilibrium to satisfy labor market clearing.

With higher consumption by high-income households and lower consumption by low-income households, consumption inequality clearly increases as the number of firms decreases. In equilibrium, some of the labor endowment of poor households, $1 - \alpha$, is used to produce goods consumed by rich households.

3.3.2 Superstar firms and misallocation

After establishing that a larger number of competitors reduces consumption inequality in the homogeneous firm case, we turn to investigate the role of dispersion in market power across firms as a driver of consumption inequality. To do so, we focus on an individual sector populated by two firms that differ in their brand values. We show that a rise in the dispersion in brand value, increasing the market power of the larger firm, results in higher consumption inequality.

Without loss of generality, let the brand value of one of the two firms be $\eta_1 = 0$. The brand value of the other firm is $\eta_2 = \Delta \geq 0$. We refer to these two firms as *small* and *large*, respectively. The comparative statics analysis considers a change in brand value dispersion in this specific sector, assuming no change in other sectors. As a result, the aggregate marginal price indices, P_h and P_l , remain constant throughout this analysis.

PROPOSITION 7. *A rise in brand value dispersion, Δ , increases consumption inequality in the sector.*

As the large firm's brand value rises relative to its competitor, its market share is larger for any given quality-price bundle. With a larger market share, it faces a lower demand elasticity and

chooses to raise prices on its products. The rise in prices, however, increases the scope for within-firm cannibalization as high-income consumers are tempted to switch to the low-end product. The large firm thus distorts its low-quality product and raises its price, while reducing the price charged for the high-quality item. Relative to the island economy, its market share among the rich increases, which undoes some of the baseline misallocation across firms in that consumer segment. Faced with lower quality for a higher price, low-income consumers increasingly purchase from the small firm, increasing misallocation there.

4 Empirical Evidence

The theory developed in the previous section delivers two main empirical predictions. First, firms with sufficient market power distort the quality of their low-end products downward, not because producing higher quality is prohibitively costly, but because doing so would cannibalize sales of their high-end offerings. Second, more concentrated markets exhibit greater dispersion in firms' market shares across income groups: dominant firms capture a disproportionate share of high-income consumers while losing low-income consumers to smaller competitors.

We present evidence for each of these predictions in turn. We first discuss examples, drawn from historical and contemporary markets, in which firms with market power have deliberately degraded the quality of low-end products to segment consumers by willingness to pay. We then turn to a systematic analysis using the NielsenIQ Consumer Panel dataset, which allows us to test whether the predicted relationship between market concentration and the income composition of firms' customer bases holds across a broad set of product categories.

4.1 Historical and contemporary examples

The mechanism at the center of our model, that firms with market power deliberately degrade the quality of their low-end products, has a long history in practice. Before turning to systematic evidence, we briefly discuss examples that illustrate this mechanism.

Railroad Transportation in 19th Century Britain. In the early decades of the British railway system, each line required its own Act of Parliament and was built and operated by a single chartered company. By the early 1840s there were over one hundred such companies, each holding a monopoly on its own route (McLean and Foster, 1992). Passengers on these lines could choose among three classes of carriage. First-class carriages were fully enclosed, padded, and relatively comfortable. Second-class carriages were roofed but sparsely upholstered. Third-class passengers, by contrast, rode in open wagons with no roof, no glazing, and wooden benches.

The conditions were not a reflection of cost considerations by railway companies given third-class passengers' willingness to pay. Dupuit (1849), a French civil engineer who observed the same practice in the French railway system, argued that the reason was instead market segmentation.⁶

⁶Our translation. Original French (p. 234): "*Ce n'est pas à cause de quelques milliers de francs qu'il en coûterait*

“It is not because of the few thousand francs it would cost to put a roof over the third-class wagons or to upholster their seats that some company or other has open wagons with wooden benches [...] one strikes at the poor, not out of a desire to make them suffer, but to frighten the rich.”

Through the lens of our model, the railway companies faced no competitive pressure on their respective routes and therefore distorted the quality of their low-end product to prevent cannibalization of first- and second-class revenue.

The British government ultimately intervened. The Railway Regulation Act of 1844, mandated that every company run at least one train per day in which third-class passengers were “protected from the weather and provided with seats,” at an average speed of at least 12 miles per hour and a fare of at most one penny per mile (McLean and Foster, 1992). Railway companies complied with the letter of the law but not its spirit. They scheduled the mandated “parliamentary trains” at the least convenient hours, used the most uncomfortable carriages the law would allow, and forced passengers to wait hours at intermediate stations. *Lloyd’s Weekly Newspaper* described these trains in 1852 as “as uncomfortable as the law will allow,” running at “a pace that wearies out all patience” (Wolmar, 2007). Lord Carlingford later admitted in Parliament that the result was “virtually a fourth class of passenger carriage” that was “in all other respects objectionable.” Even when regulation tried to impose a minimum quality floor, monopolists found margins along which to degrade service. The lesson extends beyond railways: when quality distortions are rooted in market structure, regulation may be a poor substitute for competition.

Hardware Crippleware. The same mechanism operates in contemporary technology markets, where it has acquired a name: *hardware crippleware*. The term refers to products in which a manufacturer deliberately disables features of a physically capable device to create a lower-tier version.⁷

A well-known example is Intel’s 486SX microprocessor, introduced in April 1991. At the time, Intel controlled more than 70% of the PC microprocessor market (European Commission, 2009). The 486SX was physically identical to the higher-end 486DX, but with the math coprocessor disabled. Deneckere and McAfee (1996) discuss the 486SX as a canonical example of ‘damaged goods,’ arguing that disabling the FPU is consistent with a strategy of market segmentation: by offering a chip with reduced functionality, Intel could charge a premium for the full-featured version without losing price-sensitive customers entirely. Intel’s dominance meant that cannibalization of the 486DX was a first-order concern. Disabling the coprocessor was the firm’s response.

This practice is common in markets where a small number of firms hold dominant positions. In 2001, IBM introduced “Capacity on Demand” for its zSeries enterprise servers, shipping systems with processor cores physically present but inactive at delivery. Customers could later activate this dormant capacity via software keys upon payment, without replacing or modifying the hardware. In

pour mettre un toit sur les wagons de troisième classe ou pour rembourrer les banquettes de troisième classe, que telle compagnie a des wagons découverts avec des bancs de bois [...] l’on frappe sur le pauvre, non pas qu’on ait envie de le faire souffrir personnellement, mais pour faire peur au riche.”

⁷We note that hardware crippleware is not fully captured in our model, as we assume a linear quality production cost. One way to extend our model to capture such behavior is to have a step function for the production cost.

2016, Tesla began selling the Model S 60, which contained the same 75 kWh battery pack as the Model S 75 but was software-locked to 60 kWh. Customers could unlock the full capacity for an additional fee (Weintraub, 2016). In 2020, NVIDIA introduced its Ampere-based GeForce RTX 3080 and 3090 graphics processors, which were built on the same GA102 silicon but differed in the number of enabled streaming multiprocessors and memory configuration (NVIDIA Corporation, 2020).

What unites these examples, and what connects them to our theoretical framework, is the role of market power. Quality degradation of this kind is not a generic feature of multi-product firms. It arises when a firm’s market position is strong enough that cannibalization of its high-end products becomes a binding concern. A competitive firm facing the same production technology would have no reason to disable a coprocessor or remove a roof from a railway carriage. Doing so would simply make it lose customers to rivals offering the undegraded product.

4.2 Evidence from grocery store goods

We now test the second main prediction of our model: the more concentrated a market is, the more the high-brand-value firms distort downwards the product offered to low-income households and reduce prices at the top. As a result, their market share among high-income households increases, while their popularity among low-income households declines. To see whether this pattern holds across a broad set of markets, we turn to the NielsenIQ Consumer Panel dataset. We use this particular setting, since it allows us to observe both the income of a household and detailed spending on all brands and products.

4.2.1 Data and Descriptives

We use the NielsenIQ Consumer Panel dataset from 2007–2020 to conduct our analysis. We focus on core grocery goods, which include Health & Beauty Care, Dry Grocery, Frozen Foods, Dairy, and Non-Food Grocery. The analogue to a market in our model is a product module, e.g., “olive oil” or “deodorants”. Our sample covers 869 product modules. In each of these markets, every year, we analyze the behavior of the top firm, defined as the firm with the largest market share. For robustness, we also consider the top three firms in each market. We define high- and low-income households as those with above and below median income, respectively.

Table 1 presents descriptive statistics. The average market share of the top firm is 32.4%. This firm controls, on average, 33.3% of the high-income market segment and 31.7% of the low-income market segment. Firms offer multiple products within each module. Panel B of the table shows that the median top firm offers products with 99 distinct Universal Product Codes (UPCs) across 23 different sizes.

In our model, firms offer products of different qualities. To gauge at the extent to which firms offer products of heterogeneous qualities, we proceed in two steps. First, Panel C of Table 1 shows that high-income households purchase more expensive products. On average, the products they purchase are about 10% more expensive. This difference cannot be fully explained by size variation. The

Table 1: Descriptive Statistics

	Mean	Median	S.D.
<i>Panel A: Market Structure</i>			
Top firm market share	0.324	0.299	0.167
Top firm market share, high-income	0.333	0.302	0.179
Top firm market share, low-income	0.317	0.296	0.161
HHI	0.185	0.151	0.136
<i>Panel B: Product Variety (Top Firm)</i>			
Distinct UPCs per module-year	167.2	99	251.2
Distinct sizes per module-year	35.3	23	38.7
<i>Panel C: Size and Price (Top Firm)</i>			
Relative log price (high – low income)	0.101	0.083	0.084
Relative log size (high – low income)	0.085	0.063	0.094
Relative log price per unit (high – low income)	0.016	0.009	0.059

Notes: The sample covers 869 product modules across 11,757 module-year observations. All statistics are weighted using the total sales in the module. Top firm is the firm with the largest overall market share in each module-year.

price they pay per unit of good, e.g., price per oz, is about 1.6% higher.⁸ Second, we study the variation of product attributes within firms. 96% of sales are in module-year pairs where the top firm offers products who differ by at least one product attributes.⁹ In particular, 72.4% of sales are in module-years where the top firm offers both an organic and non-organic versions of their product.

4.2.2 Market concentration and relative market shares

Our key variable of interest is the relative market share of the top firm, defined as the top firm’s market share among high-income households divided by its market share among low-income households. In the absence of cannibalization concerns within the firm, the relative market share would be one in our model; firms are equally popular across all income segments. The lower the external competition a firm faces, the more it tilts qualities and prices in favor of high-income consumers, increasing its market share in that segment while losing low-income customers. We therefore test whether the relative market share of the top firm varies systematically with market concentration. Our baseline measure of market concentration within each module is the Herfindahl–Hirschman Index (HHI). Our baseline regression takes the following form:

$$\log \left(\frac{s_{mt}^H}{s_{mt}^L} \right) = \delta_t + \beta \log HHI_{mt} + \epsilon_{mt}, \quad (4.1)$$

where s_{mt}^H and s_{mt}^L are the market share of the top firm in market m in year t in the high- and low-

⁸Quantity discounts are prevalent in this market, see [Bornstein and Peter \(2025\)](#). This feature of the data would lead to lower price-per-unit for larger sizes, indicating that high-income households indeed purchase higher quality goods.

⁹We consider the following product attributes: (i) organic claim, (ii) type, (iii) flavor, (iv) form, (v) formula, (vi) style, and (vii) scent.

income market segments, respectively. We control for time fixed effects, δ_t , and study the elasticity of the relative market share with respect to the HHI in the market.

Table 2 reports the results. Column (1) presents the benchmark specification. The elasticity of relative market shares with respect to HHI is 0.074 and is statistically significant. To interpret the economic magnitude, note that the median relative log market share is 0.042. The coefficient implies that a one standard deviation increase in log-HHI raises it from 0.042 to 0.096. Column (2) shows that the same pattern holds if we simply compare markets with above median HHI to ones below it. The difference is 8.9 log points.

Through the lens of our model, the distortions the firm introduces directly depend on its own market share. Specification (3) in Table 2 considers the same regression specification, replacing the market HHI with the firm’s own market share. This specification also yields a positive and significant coefficient. The elasticity of the relative market share with respect to the firm’s own market share is 0.096. That is, the relative market share of a firm with a 10% larger market share is about 1% higher.¹⁰

Table 2: Concentration and relative market share

	Top Firm			Top-3 Firms		
	(1)	(2)	(3)	(4)	(5)	(6)
log(HHI)	0.074*** (0.021)			0.037*** (0.011)		
1 [HHI > median]		0.089*** (0.028)			0.046*** (0.014)	
log(own market share)			0.096*** (0.025)			0.073*** (0.023)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	11,757	11,757	11,757	11,757	11,757	11,757
Clusters	869	869	869	869	869	869
R^2	0.089	0.069	0.087	0.116	0.098	0.117

Notes: The dependent variable is $\log(s_{f,m,t}^H/s_{f,m,t}^L)$, where s^H and s^L denote the market share of firm f among high- and low-income households in module m and year t . Columns (1)–(3) use the top firm by overall market share; columns (4)–(6) use the top three firms combined. All regressions include year fixed effects and are weighted by module-year sales. Standard errors clustered by product module are in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Finally, columns (4–6) repeat the analysis but consider the combined market share of the top three firms instead of only the top firm. The coefficients are all positive and significant, further strengthening the hypothesis that lack of external competition tilts the market share of top firms towards the high-income segment. Moreover, consistent with the model’s prediction, these coefficients are all smaller. The firms ranked 2–3 do not have as much market power as the top firm, and would therefore not be expected to distort their product portfolios as much as the top firm.

¹⁰Note this is not 10 percentage points higher market share, but 10 percent. i.e., a firm with a market share of 33% relative to a firm with a market share of 30%.

5 Quantitative Results

To gauge how much low-income households are losing from low competition across firms, we calibrate the model to the 2022 US economy. We find that consumption inequality is 7% higher than in the efficient allocation. That is, the presence of market power substantially amplifies inequality. We then calibrate the model to the lower level of market concentration on the 1980s and quantify the unequal incidence of the rise in firms’ market power since.

5.1 Calibration

We assume that the brand value of firms is drawn from a Pareto distribution with shape χ . Without loss of generality, we normalize the scale of the Pareto distribution to 1. We maintain the assumption that there are two types of households who differ by their income. The model has 4 parameters to be calibrated: (i) the degree of inequality α , (ii) the number of firms N , (iii) the shape of the Pareto distribution χ , and (iv) the dispersion of the idiosyncratic taste shocks, σ . We calibrate σ , which governs the elasticity of substitution across brands, to 3 following [Midrigan \(2011\)](#).

Table 3 lists the calibrated parameter values as well as the data moments used to identify them. The degree of income inequality, α , is calibrated using evidence on expenditure inequality from the Panel Study of Income Dynamics (PSID). The resulting value of α is 0.31. The parameter is chosen to match the relative expenditures of households above and below median income. We match expenditure inequality rather than income inequality as our model is static and abstracts from saving.

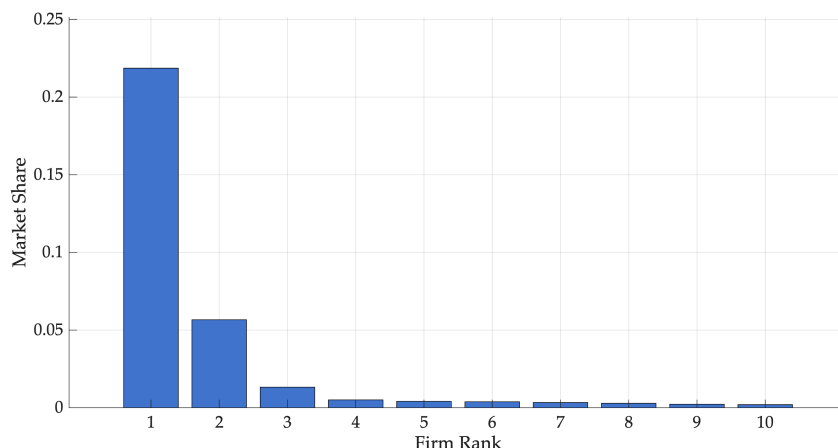
We calibrate the market structure parameters, N and σ , to match the average market structure of a 6-digit NAICS industry using the Firm Size Statistics data tables from the 2022 Economic Census. In particular, we set N to 7,000 to equal the average number of firms within each industry. And we calibrate the Pareto shape parameter χ , which controls the brand value dispersion across firms, to match an average HHI of 0.05, resulting in a shape distribution of 7.3. The large dispersion in brand value across firms allows the model to match the degree of concentration observed in the data despite the large number of operating firms.

Table 3: Calibrated Parameters

Parameter	Description	Value	Moment	Source
σ	Dispersion of taste shocks	3	EoS across firms	Literature
α	Spending inequality	0.31	Non-housing exp. inequality	PSID (2022)
N	Number of firms	7K	Avg. # firms in 6-digit NAICS	Census (2022)
χ	Pareto shape	7.3	Avg. HHI 6-digit NAICS (0.05)	Census (2022)

Notes: his table reports the calibrated parameters, the target moments, as well as their model equivalent.

Figure 1: Market shares of top 10 firms



Notes: The figure presents the distribution of market shares across the top ten firms in the calibrated economy. The top firm controls 22% of the market, and the market share drops exponentially as we go down the brand-value ladder. This exponential decline is governed by the Pareto shape χ and the elasticity of substitution σ , and is calibrated to match the degree of HHI observed in the data.

5.2 Consumption Inequality

Let us now turn to study the degree of consumption inequality in the calibrated economy. Recall that absent cannibalization considerations, the degree of consumption inequality mirrors that of income inequality. Cannibalization considerations lead high-market-power firms to distort down the quality of their low-end products and to raise their price. This force increases consumption inequality beyond income inequality.

Quantitatively, we find the consumption inequality is 6.8 percentage points larger than income inequality. Table 5 presents our findings. Relative to the efficient allocation, low-income households' welfare is 4.8% lower in the market allocation, while high-income households are 2% better off.

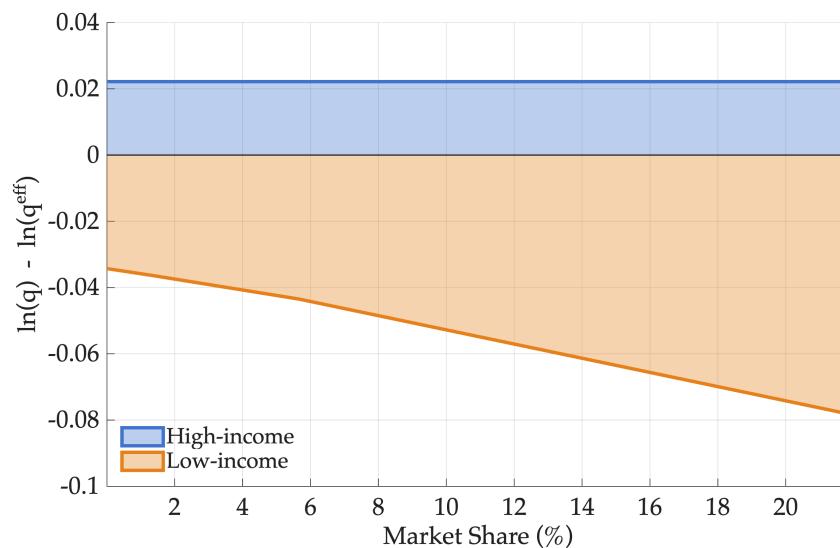
Table 4: Aggregate Consumption relative to the Efficient Allocation

	High-income	Low-income
Market equilibrium	+2.0%	-4.8%
Only quality distortions	+2.2%	-4.6%
Only misallocation	-0.2%	-0.4%
Island economy	-0.3%	-0.3%

Notes: This table reports the log difference in aggregate consumption of each income type between the various allocations and the efficient allocation. The first row corresponds to the market allocation; the second row considers only market distortion in quality offered, not in market shares; the third row keeps qualities at the efficient level and only considers distortions in market shares; and the last column presents results for the island economy.

The higher consumption inequality in the market equilibrium is driven by both quality distortions and misallocation. On the quality margin, high-brand value firms distort the quality level of their low-end products downwards. Figure 2 illustrates this. Relative to the efficient allocation, the quality level offered to low-income households lower at all firms (orange line). For lower-brand-value firms,

Figure 2: Quality distortions relative to efficient allocation



Notes: This figure plots the quality each firm offers relative to what it would be in the efficient allocation. Numbers are reported as log differences. The blue line corresponds to high-income households, the orange one to low-income households. The x-axis measures the firms overall market share across both segments combined.

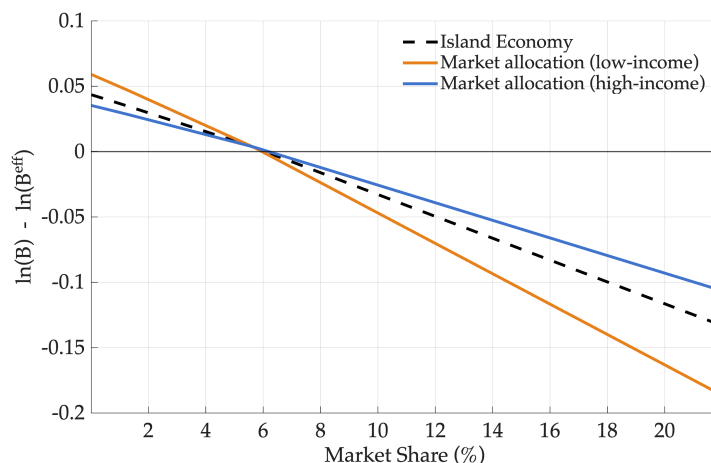
cannibalization concerns are less pronounced and the distortion is around 3%. It reaches as high as 8% at high brand-value firms. On the flip side, high-income consumers are being offered 2.2% higher quality across all firms (blue line).

The distortion in quality levels is responsible for nearly all the additional consumption inequality present in the market allocation. In the second row of Table 5, we report the aggregate consumption of both income groups if there was only distortions in qualities but the allocation of consumers across firms was equal to the efficient one. The 2.2% higher quality directly translates into 2.2% higher aggregate welfare for the rich, while low-income households, whose quality is distorted downwards by 3-8%, are 4.6% worse off in the aggregate.

The second channel—misallocation of consumers across firms—plays a smaller role. In the island economy, where there are no quality distortions, aggregate consumption of all households is 0.3% lower than in the efficient allocation. If there was only the cross-firm misallocation channel, consumption would be 0.2% lower for the rich and 0.4% lower for the poor. That is, misallocation across firms worsens for low-income households, who are increasingly driven away by high-brand-value firms and improves for high-income ones.

Figure 3 plots the market shares of each firm relative to the efficient allocation in the island economy as well as in the market economy for both income groups. All three lines are different from zero and downward sloping, meaning that there is misallocation of consumers across firms. High-brand-value firms, i.e., those with larger overall market shares, have too few consumers, while too many households buy from low-brand-value firms that charge lower markups. The allocation of consumers across firms improves for high-income households in the market allocation: the blue line is flatter than the island economy. For low-income households, distortions in quality and higher prices at top firms exacerbate

Figure 3: market shares relative to efficient allocation



Notes: This figure plots each firm’s market share relative to the efficient allocation. Numbers are reported as log differences. The dashed line correspond to the island economy, in which market shares are identical across the consumer segments. The solid lines correspond to the market allocation: high-income households in blue, low-income ones in orange. The x-axis measures the firms overall market in the market allocation.

the misallocation present in the island economy.

5.3 Distributional impact of the rise of market power

Since the 1980s, there has been a secular rise in market concentration in the U.S. To gauge whether this rise differentially affected households of different income groups, we re-calibrate the model to match the level of concentration 40 years ago.

Smith and Ocampo (2025) estimate that the average HHI in US retail was 0.02 in 1980, compared to 0.05 in 2022. To match the lower level of concentration, we re-estimate the tail index of the Pareto distribution. Conceptually, we attribute the rise in concentration to the emergence of highly productive “super-star firms” (Autor et al., 2020). To match the lower HHI, we calibrate a less dispersed Pareto distribution of brand values (tail index of 7.75 instead of 7.3). This implies that the top 4 firms in the market account for 16% of sales rather than close to 30% in 2022.¹¹

In the 1980s, the top firms were less dominant in their market and consequently yielded less market power. With lower markups across the board, high-income households were less tempted to purchase low-end products and firms had less incentives to distort their product portfolios. All in all, consumption inequality between rich and poor was only 5.9%, nearly 1 percentage point lower than in 2022.

¹¹We assume that the scale of the Pareto distribution remains unchanged. This implies that the average brand value is lower than in our baseline calibration. Note that the average brand value affects aggregate welfare, but not inequality between rich and poor.

Table 5: The Distributional Impact of Rising Market Power

	1980s	2022
<i>Calibration</i>		
Average HHI	0.02	0.05
Pareto tail index	7.75	7.3
<i>Results</i>		
Excess consumption inequality	5.9%	6.8%

Notes: This table presents the calibration and results for our baseline calibration (2022 column) and the recalibrated version targeting HHI in the 1980s.

6 Conclusion

This paper studies the distributional consequences of market power through the lens of endogenous product design. We show that firms with sufficient market power distort the quality of their low-end products to prevent cannibalization of their high-end offerings, a force that is absent under sufficient competition. Through general equilibrium forces, these distortions benefit high-income households at the expense of low-income ones. Both a reduction in the number of competitors and the emergence of superstar firms amplify consumption inequality through this channel.

Empirical evidence from the NielsenIQ Consumer Panel supports the model’s key prediction: in more concentrated markets, dominant firms capture a disproportionately larger share of high-income consumers. In a calibrated version of the model, the distributional costs of market power are large. Market power reduces the welfare of low-income households by 4.8% relative to the efficient allocation while increasing that of high-income households by 2.0%, effects that are an order of magnitude larger than the standard efficiency costs of imperfect competition. The rise in market concentration since 1980 widened the welfare gap between rich and poor by 0.9 percentage points.

Several questions remain open. Our analysis takes market structure as given: the number of firms and their brand values are exogenous. In practice, both may respond to the forces we study. The distributional consequences of market power could feed back into concentration through entry and exit. And if a firm’s brand value depends in part on the composition of its customer base, the mechanism we identify may be self-reinforcing: by distorting their low-end products and tilting their customer base toward high-income households, dominant firms may further increase their brand appeal, strengthening the very market power that led to the distortion. Whether promoting competition is more effective at protecting low-income consumers than direct quality regulation, which the historical record suggests is difficult to enforce, is a question we leave for future research.

References

- ATKESON, A. AND A. BURSTEIN (2008): “Pricing-to-market, trade costs, and international relative prices,” *American Economic Review*, 98, 1998–2031. 1, 3.1
- AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (2020): “The Fall of the Labor Share and the Rise of Superstar Firms*,” *The Quarterly Journal of Economics*, 135, 645–709. 1, 1, 5.3
- BAQAEI, D. R. AND E. FARHI (2020): “Productivity and misallocation in general equilibrium,” *The Quarterly Journal of Economics*, 135, 105–163. 1
- BECKER, J. (2024): “Do Poor Households Pay Higher Markups in Recessions?” Working paper, New York University. 1
- BOAR, C. AND V. MIDRIGAN (2024): “Markups and Inequality,” *Review of Economic Studies*, advance online publication. 1
- BORNSTEIN, G. AND A. PETER (2025): “Nonlinear Pricing and Misallocation,” *American Economic Review*, 115, 3853–3908. 8
- DE LOECKER, J., J. ECKHOUT, AND G. UNGER (2020): “The rise of market power and the macroeconomic implications,” *The Quarterly Journal of Economics*, 135, 561–644. 1, 1
- DENECKERE, R. J. AND R. P. MCAFEE (1996): “Damaged Goods,” *Journal of Economics & Management Strategy*, 5, 149–174. 4.1
- DUPUIT, J. (1849): “De l’influence des péages sur l’utilité des voies de communication,” *Annales des Ponts et Chaussées*, 17, 170–248, english translation in *International Economic Papers*, 11:7–31, 1962. 4.1
- EDMOND, C., V. MIDRIGAN, AND D. Y. XU (2022): “How costly are markups?” *Journal of Political Economy*. 1
- EUROPEAN COMMISSION (2009): “Commission Decision of 13 May 2009, Case COMP/C-3/37.990 — Intel,” OJ C 227, 22.9.2009, pp. 13–17. 4.1
- FINKELSTEIN, M. O. AND R. M. FRIEDBERG (1967): “The application of an entropy theory of concentration to the Clayton Act,” *The Yale Law Journal*, 76, 677–717. 2.3
- HSIEH, C.-T. AND P. J. KLENOW (2009): “Misallocation and manufacturing TFP in China and India,” *The Quarterly journal of economics*, 124, 1403–1448. 1
- JARAVEL, X. (2019): “The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector,” *The Quarterly Journal of Economics*, 134, 715–783. 1
- KAPLAN, G. AND G. MENZIO (2015): “The Morphology of Price Dispersion,” *International Economic Review*, 56, 1165–1206. 1

- KARABARBOUNIS, L. AND B. NEIMAN (2014): “The global decline of the labor share,” *The Quarterly journal of economics*, 129, 61–103. 1
- MCLEAN, I. AND C. D. FOSTER (1992): “The Political Economy of Regulation: Interests, Ideology, Voters and the UK Regulation of Railways Act 1844,” *Public Administration*, 70, 313–331. 4.1
- MIDRIGAN, V. (2011): “Menu costs, multiproduct firms, and aggregate fluctuations,” *Econometrica*, 79, 1139–1180. 5.1
- MONGEY, S. AND M. E. WAUGH (2025): “Pricing Inequality,” Working Paper 33399, NBER. 1
- MUSSA, M. AND S. ROSEN (1978): “Monopoly and product quality,” *Journal of Economic Theory*, 18, 301–317. 1, 1
- NORD, L. (2023): “Shopping, Demand Composition, and Equilibrium Prices,” . 1
- NVIDIA CORPORATION (2020): “NVIDIA Ampere GA102 GPU Architecture Whitepaper,” Technical whitepaper, NVIDIA Corporation, accessed March 2026. 4.1
- OBERFIELD, E. (2023): “Inequality and Measured Growth,” NBER Working Paper 31096, National Bureau of Economic Research. 1
- ROCHET, J.-C. AND L. A. STOLE (2002): “Nonlinear Pricing with Random Participation,” *The Review of Economic Studies*, 69, 277–311. 1
- SANGANI, K. (2024): “Markups across the income distribution: Measurement and implications,” Available at SSRN 4092068. 1
- SMITH, D. A. AND S. OCAMPO (2025): “The evolution of US retail concentration,” *American Economic Journal: Macroeconomics*, 17, 71–101. 5.3
- STOLE, L. A. (1995): “Nonlinear Pricing and Oligopoly,” *Journal of Economics & Management Strategy*, 4, 529–562. 1
- WEINTRAUB, S. (2016): “Tesla’s New 60kWh Pricing Option Is a Software Revolution,” *Electrek*. 4.1
- WOLMAR, C. (2007): *Fire and Steam: A New History of the Railways in Britain*, London: Atlantic Books. 4.1

A Mathematical appendix

A.1 Proof of Lemma 1

In equilibrium, each household chooses the product quality that is designated to them. The discrete choice problem of the household is¹²

$$\max_j = \arg \max_j (\sigma - 1) \left(u(q_{jl}) - \frac{p_{jl}}{P_l} + \eta_j \right) + \epsilon_{ij}. \quad (\text{A.1})$$

Let $s_{ij} \equiv (\sigma - 1)(u(q_{jl}) - \frac{p_{jl}}{P_l} + \eta_j)$ denote the surplus of consumer i when purchasing their designated bundle from firm j , net of the taste shock ϵ_{ij} . That is, s_{ij} is identical for all low-income (s_{jl}) and all high-income (s_{jh}) consumers.

The probability that firm j is chosen, conditional on the taste shock of the household being $\bar{\epsilon}$ is given by

$$\begin{aligned} Pr(j_l^* = j | \epsilon_{ij} = \bar{\epsilon}) &= \prod_{j' \neq j} [Pr(s_{jl} + \bar{\epsilon} \geq s_{j'l} + \epsilon_{ij'})] \\ &= \prod_{j' \neq j} [Pr(\epsilon_{ij'} \leq s_{jl} - s_{j'l} + \bar{\epsilon})] \\ &= \prod_{j' \neq j} e^{-e^{-(s_{jl} - s_{j'l} + \bar{\epsilon})}} \\ &= \prod_{j'} \left[e^{-e^{-(s_{jl} - s_{j'l} + \bar{\epsilon})}} \right] e^{e^{-\bar{\epsilon}}} \\ &= e^{-\sum_{j'} e^{-(s_{jl} - s_{j'l} + \bar{\epsilon})}} e^{e^{-\bar{\epsilon}}} \\ &= e^{-e^{-\bar{\epsilon}} \sum_{j'} e^{-(s_{jl} - s_{j'l})}} e^{e^{-\bar{\epsilon}}} \\ &= e^{-e^{-\bar{\epsilon}} Q} e^{e^{-\bar{\epsilon}}}, \end{aligned}$$

where $Q = \sum_{j'} e^{-(s_{jl} - s_{j'l})}$. The market share of firm j for low-income households is derived by integrating over the probability distribution of ϵ_{ij} :

$$\begin{aligned} B_{jl} &= \int e^{-Qe^{-\bar{\epsilon}}} e^{e^{-\bar{\epsilon}}} e^{-(\bar{\epsilon} - \mu) - e^{-\bar{\epsilon}}} d\bar{\epsilon} \\ &= \int e^{-Qe^{-\bar{\epsilon}}} e^{-\bar{\epsilon}} d\bar{\epsilon} \\ &= \int_0^\infty e^{-Qx} dx \\ &= \frac{1}{Q} \\ &= \frac{e^{s_j}}{\sum_{j'} e^{s_j}}. \end{aligned}$$

¹²We have multiplied the maximand by $\sigma - 1$. This does not change the argmax. It will make the mathematical derivations simpler as will be clear below.

where the third equality uses integration by substitution with $x = e^{-\epsilon}$. That is,

$$B_{jl} = \frac{e^{s_j}}{\sum_{j'} e^{s_{j'}}}. \quad (\text{A.2})$$

Using the same derivations, we can obtain

$$B_{jh} = \frac{e^{s_{jl}}}{\sum_{j'} e^{s_{j'l}}}. \quad (\text{A.3})$$

That is, we have proved that

$$\mathcal{B}(q_j, p_j, \eta_j; \{q_{j'}, p_{j'}, \eta_{j'}\}_{j' \neq j}, P) = \frac{e^{(\sigma-1)(u(q_j) - \frac{p_j}{P} + \eta_j)}}{\sum_{j'} e^{(\sigma-1)(u(q_{j'}) - \frac{p_{j'}}{P} + \eta_{j'})}}. \quad (\text{A.4})$$

A.2 Efficient allocation

PROOF OF PROPOSITION 2. We solve the planner's problem under the conjecture that the incentive compatibility constraints are slack and confirm part (i) of the proposition at the end. Without the incentive compatibility constraints, the planner's problem can be written as

$$\begin{aligned} \max_{\{j_{i\tau}(s), q_{i\tau}(s)\}_{\{i,\tau,s\}}} & \int_i \int_s \left[\ln(q_{il}(s)) + \eta_{j_{il}(s)} + \frac{\varepsilon_{ij_{il}(s)}}{\sigma-1} + \ln(q_{ih}(s)) + \eta_{j_{ih}(s)} + \frac{\varepsilon_{ij_{ih}(s)}}{\sigma-1} \right] ds \\ \text{s.t.} & \int_i \int_s \kappa q_{i\tau}(s) ds = L_\tau. \end{aligned} \quad (\text{A.5})$$

(ii) Taking first-order conditions with respect to $(q_{i\tau}(s))$ directly yields

$$\frac{1}{q_{i\tau}(s)} = \frac{\kappa}{P_\tau^{\text{FB}}}, \quad (\text{A.6})$$

where $1/P_\tau^{\text{FB}}$ are the two Lagrange multipliers on the resource constraints of low-income and high-income people.

(iii) Plugging $\frac{1}{q_{i\tau}(s)} = \frac{\kappa}{P_\tau^{\text{FB}}}$ into the resource constraints directly yields that $P_\tau^{\text{FB}} = L_\tau$

(iv) The choice of $j_{i\tau}(s)$ only affects utility and has no consequences for the resource constraint. Hence, the planner chooses $j_{i\tau}(s)$ to point-wise maximize the integral.

(v) Plugging (A.6) into (A.4) yields the expression for the efficient market share

(i) Given part (iv) of Proposition 2, the social planner maximizes the choice of firm j given the consumer's idiosyncratic taste shock ε_{ij} . Hence truthfully reporting one's taste shock type is trivially optimal. Note that the planner can easily prevent low-income consumers from reporting that they have high labor endowment by making labor supply part of the allocation.

■

A.3 Microfounding Aggregate Price Index

Suppose you currently purchase from firm j , which gives utility

$$u(q_j) + \eta_j + \frac{\epsilon_{ij}}{\sigma - 1} \quad (\text{A.7})$$

and you consider instead buying from firm j'

$$u(q_{j'}) + \eta_{j'} + \frac{\epsilon_{ij'}}{\sigma - 1} \quad (\text{A.8})$$

If we define $1/P$ as the utility gain from an extra dollar, then the utility gain from an extra $p_{j'} - p_j$ dollars is

$$\frac{p_{j'} - p_j}{P} \quad (\text{A.9})$$

Given that there is a continuum of industries, there exists, for *any* aggregate P a pair of firms j and j' such that

$$\frac{1}{P} \equiv \frac{u(q_{j'}) - u(q_j) + \eta_{j'} - \eta_j + \frac{\epsilon_{ij'} - \epsilon_{ij}}{\sigma - 1}}{p_{j'} - p_j} \quad (\text{A.10})$$

A.4 Island Economy

Setup and Firm Problem In this section, we set up the island economy and derive all results discussed in the main text. Since firms can offer each type of consumer only one bundle, the incentive compatibility constraints of equation (2.5) need not hold. Further, since production is linear, the firm's problem can be fully separated across islands.

$$\begin{aligned} \pi_{j\tau} = & \max_{\{q_{j\tau}, p_{j\tau}, B_{j\tau}\}} B_{j\tau} (p_{j\tau} - \kappa q_{j\tau}), \\ \text{s.t. } & B_{j\tau} = \frac{e^{(\sigma-1)(\ln(q_{j\tau}) - \frac{p_{j\tau}}{P_\tau} + \eta_j)}}{\sum_{j'} e^{(\sigma-1)(\ln(q_{j'\tau}) - \frac{p_{j'\tau}}{P_\tau} + \eta_{j'})}}. \end{aligned} \quad (\text{A.11})$$

Taking first-order conditions, we obtain

$$\begin{aligned}
[p_{j\tau}] : \quad & B_{j\tau} + (p_{j\tau} - \kappa q_{j\tau}) \frac{\partial B_{j\tau}}{\partial p_{j\tau}} = 0, \\
[q_{j\tau}] : \quad & -\kappa B_{j\tau} + (p_{j\tau} - \kappa q_{j\tau}) \frac{\partial B_{j\tau}}{\partial q_{j\tau}} = 0.
\end{aligned}$$

Differentiating $B_{j\tau}$ and combining the two equations, we get the following optimality conditions for prices and quantities:

$$\frac{1}{q_{j\tau}} = \frac{\kappa}{P_\tau} \tag{A.12}$$

$$p_{j\tau} = \kappa q_{j\tau} + \frac{P_\tau}{(\sigma - 1)(1 - B_{j\tau})} \tag{A.13}$$

Allocations in Equilibrium First, note that the marginal price index in the island economy is identical to the first-best, $P_i = P_i^{FB}$. Since all households consume from exactly one firm in each sector, and production costs are homogeneous,

A.5 Market Allocation

Let us start by deriving the optimality conditions in the decentralized equilibrium. We do so in two steps. First, we solve the optimal allocation separately across income groups, disregarding the incentive compatibility constraint. Second, if the unconstrained allocation violates the incentive compatibility constraint, we solve the combined problem while imposing the incentive compatibility constraint holds with equality.

Unconstrained case. For ease of notation, let us drop both the consumer type and the firm indicators. The unconstrained firm's problem is given by

$$\begin{aligned}
\max_{\{q,p,B\}} \quad & B(p - \kappa q), \\
\text{s.t.} \quad & B = \frac{e^{(\sigma-1)(\ln(q) - \frac{p}{P} + \eta)}}{e^{(\sigma-1)(\ln(q) - \frac{p}{P} + \eta)} + S_{-j}},
\end{aligned} \tag{A.14}$$

Taking FOC:

$$\begin{aligned}
[B] : \quad & p - \kappa q = \lambda, \\
[q] : \quad & \kappa q = \lambda(\sigma - 1)(1 - B), \\
[p] : \quad & P = \lambda(\sigma - 1)(1 - B).
\end{aligned}$$

The last two equations imply

$$\kappa q = P.$$

Then, combining the first and third optimality conditions, we have

$$p = \kappa q + \frac{P}{(\sigma - 1)(1 - B)}. \quad (\text{A.15})$$

And plugging into the market share function:

$$B e^{\frac{1}{1-B}} = \frac{e^{(\sigma-1)(\ln P - \ln \kappa - 1 + \eta)}}{e^{(\sigma-1)(\ln P - \ln \kappa - 1 + \eta) - \frac{1}{1-B}} + S_{-j}}. \quad (\text{A.16})$$

So, in the unconstrained allocation, we have

$$\ln q_h - \ln q_l = \ln P_h - \ln P_l. \quad (\text{A.17})$$

Moreover,

$$p_h - p_l = P_h - P_l + \frac{P_h}{\sigma - 1} \frac{1}{1 - B_{jh}} - \frac{P_l}{\sigma - 1} \frac{1}{1 - B_{jl}} \quad (\text{A.18})$$

Recall, the IC constraint is

$$\ln(q_h) - \ln(q_l) \geq \frac{p_h - p_l}{P_h}, \quad (\text{A.19})$$

Plugging in the unconstrained allocation, the IC boils down to

$$-\ln\left(\frac{P_l}{P_h}\right) \geq 1 + \frac{1}{\sigma - 1} \frac{1}{1 - B_{jh}} - \frac{P_l}{P_h} \left(1 + \frac{1}{\sigma - 1} \frac{1}{1 - B_{jl}}\right), \quad (\text{A.20})$$

Constrained allocation. Suppose that the IC-high constraint is binding, and let us characterize the constrained allocation. The firm's problem is given by

$$\begin{aligned} & \max_{\{q_l, q_h, p_l, p_h, B_l, B_h\}} B_l (p_l - \kappa q_l) + B_h (p_h - \kappa q_h), & (\text{A.21}) \\ & \text{s.t.} \quad B_l = \frac{e^{(\sigma-1)\left(\ln(q_l) - \frac{p_l}{P_l} + \eta_j\right)}}{e^{(\sigma-1)\left(\ln(q_l) - \frac{p_l}{P_l} + \eta_j\right)} + S_{-jl}}, \\ & \quad B_h = \frac{e^{(\sigma-1)\left(\ln(q_h) - \frac{p_h}{P_h} + \eta_j\right)}}{e^{(\sigma-1)\left(\ln(q_h) - \frac{p_h}{P_h} + \eta_j\right)} + S_{-jh}}, \\ & \quad \ln(q_h) - \ln(q_l) = \frac{p_h - p_l}{P_h} \end{aligned}$$

Taking FOC:

$$\begin{aligned}
[B_l] : \quad & p_l - \kappa q_l = \lambda_l, \\
[B_h] : \quad & p_h - \kappa q_h = \lambda_h, \\
[q_l] : \quad & \kappa q_l = \lambda_l(\sigma - 1)(1 - B_l) - \psi \frac{1}{B_l}, \\
[q_h] : \quad & \kappa q_h = \lambda_h(\sigma - 1)(1 - B_h) + \psi \frac{1}{B_h}, \\
[p_l] : \quad & 1 = \lambda_l(\sigma - 1)(1 - B_l) \frac{1}{P_l} - \psi \frac{1}{P_h B_l}, \\
[p_h] : \quad & 1 = \lambda_h(\sigma - 1)(1 - B_h) \frac{1}{P_h} + \psi \frac{1}{P_h B_h},
\end{aligned}$$

where λ_h and λ_l are the Lagrange multipliers on the market share demand functions and ψ is the Lagrange multiplier on the IC constraint. Combining the FOC wrt q_h and p_h , we get that the quantity sold to the high-type equates marginal utility to marginal cost

$$q_h = \frac{P_h}{\kappa}. \quad (\text{A.22})$$

So that $\lambda_h = p_h - P_h$ from the second optimality condition. Combining the the optimality conditions w.r.t. q_l and p_l , we obtain

$$q_l = \frac{P_l}{\kappa} - \frac{\psi}{\kappa B_l} \left(1 - \frac{P_l}{P_h}\right). \quad (\text{A.23})$$

The equation above shows that the IC constraint leads to a downward distortion in the low-income quality bundle ($q_l < \frac{P_l}{\kappa}$). The optimality conditions w.r.t. q_h and B_h yield

$$\psi = B_h P_h - (p_h - P_h)(\sigma - 1)B_h(1 - B_h). \quad (\text{A.24})$$

Finally, the optimality conditions w.r.t. q_l and B_l imply:

$$\kappa q_l = (p_l - \kappa q_l)(\sigma - 1)(1 - B_l) - \psi \frac{1}{B_l}. \quad (\text{A.25})$$

PROOF OF PROPOSITION 6.

Since all firms are homogeneous, we have that $B_{j\tau} = \frac{1}{N}$ for all firms and all consumer types.

The proof proceeds in 3 steps

1. If the IC is binding, consumption inequality is higher than income inequality and increasing in N .
2. If the IC is slack, consumption inequality is equal to income inequality.
3. If the IC is slack for N , it is slack for all $N' > N$.

Part 1 When firms are identical, they all sell the same quantity q_i for the same price p_i to each consumer type $i \in l, h$. The budget constraints of the two types of agents thus become

$$p_h = (1 + \alpha)(1 + \pi) \quad (\text{A.26})$$

$$p_l = (1 - \alpha)(1 + \pi) \quad (\text{A.27})$$

The resource constraint implies that the total quantity of labor, 1 must be used for production of either the high or the low bundle,

$$\kappa(q_l + q_h) = 1 \quad (\text{A.28})$$

Let $\tilde{\sigma}_N \equiv (\sigma - 1)(1 - 1/N)$

From the FOCs of the firm's problem wrt B_i and q_i , we have that

$$q_l = \frac{\tilde{\sigma}_N p_l - N\psi}{\kappa(1 + \tilde{\sigma}_N)} \quad (\text{A.29})$$

$$q_h = \frac{\tilde{\sigma}_N p_h + N\psi}{\kappa(1 + \tilde{\sigma}_N)} \quad (\text{A.30})$$

Combining this with the resource constraint (A.28), we get a closed-form expression for profits and consequently the prices paid by low and high types.

$$\pi = \frac{1}{\tilde{\sigma}_N} \quad (\text{A.31})$$

$$p_h = (1 + \alpha) \left(1 + \frac{1}{\tilde{\sigma}_N} \right) \quad (\text{A.32})$$

$$p_l = (1 - \alpha) \left(1 + \frac{1}{\tilde{\sigma}_N} \right) \quad (\text{A.33})$$

Now let's show that $\ln\left(\frac{q_h}{q_l}\right)$ is a decreasing function of N . That is, when N decreases, consumption inequality increases. Let $L \equiv \ln\left(\frac{q_h}{q_l}\right)$ be the object whose properties we want to prove. Further, let $s(L) \equiv \frac{e^L}{1+e^L}$. Using the definition of L , $s(L) = \frac{q_h}{q_l + q_h}$.

Recall that the resource constraint is $q_l + q_h = 1/\kappa$. Re-arranging and multiplying by q_h , and using the optimality condition for q_h , $P_h/\kappa = q_h$, we have

$$\frac{q_h}{q_l + q_h} = \kappa q_h \quad (\text{A.34})$$

$$q_h = s(L) \frac{1}{\kappa} \quad (\text{A.35})$$

$$P_h = s(L) \quad (\text{A.36})$$

Using the incentive compatibility

$$L = \frac{p_h - p_l}{P_h} \quad (\text{A.37})$$

$$L = \frac{\alpha \left(1 + \frac{1}{\sigma N}\right)}{s(L)} \quad (\text{A.38})$$

$$s(L)L = \alpha \left(1 + \frac{1}{(\sigma - 1)(1 + 1/N)}\right) \quad (\text{A.39})$$

Equation (A.39) is a single equation that relates relative log consumption to parameters of the model. Since $s(L)$ is an increasing function of L , the LHS of (A.39) is increasing. The RHS of (A.39) is decreasing in N , finalizing the proof.

Part 2 When the IC is slack, the market allocation is equal to the island economy. By Proposition 3, consumption inequality equals income inequality

Part 3 Using the fact that $B_{j\tau} = \frac{1}{N}$, we can re-write the incentive compatibility constraint in (A.20) as

$$\frac{-\ln(P_l/P_h)}{1 - P_l/P_h} \geq 1 + \frac{1}{\sigma - 1} \frac{N}{N - 1} \quad (\text{A.40})$$

Suppose the IC is slack for N . With a slack IC, the allocations are identical to the ones in the island economy, and $P_l/P_h = (1 - \alpha)/(1 + \alpha)$. That is, as long as the IC is slack, the LHS of (A.40) is independent of N .

The RHS of (A.40) is decreasing in N . Hence, if the IC is slack for N , it is slack for N' . Since the IC is slack, P_l/P_h is indeed equal to $(1 - \alpha)/(1 + \alpha)$.

■