The Innovation Engel Curve

Vasco Carvalho* Nir Jaimovich† Sergio Rebelo‡

September 14, 2026

Abstract

This paper studies the impact of income inequality on the composition of consumption demand and long-run economic growth. Using a unique large-scale U.S. dataset, we show that the consumption bundle of the top 1 percent contains a markedly smaller share of goods produced by innovation-intensive firms than that of the remaining 99 percent of households. This finding implies that reallocating income toward the top 1 percent reduces aggregate demand for innovation-intensive goods, weakens firms' incentives to create new varieties, and ultimately slows economic growth.

JEL Classification: D11, D31, O31, O41.

Keywords: Income inequality; product innovation; nonhomothetic demand; market size; mass consumption;

*University of Cambridge and CEPR
†University of California, San Diego, NBER and CEPR
‡Northwestern University, NBER and CEPR

1 Introduction

In 1956, Walter Reuther, president of the United Automobile Workers, recounted the following episode:[1] "I went through this Ford engine plant about three years ago, when they first opened it. There are acres and acres of machines, and here and there you will find a worker standing at a master switchboard, just watching, green and yellow lights blinking off and on, which tell the worker what is happening in the machine. One of the management people, with a slightly gleeful tone in his voice said to me, "How are you going to collect union dues from all these machines?" And I replied, "You know, that is not what's bothering me. I'm troubled by the problem of how to sell automobiles to these machines.""

Reuther's question is likely to become increasingly relevant in an economy shaped by artificial intelligence (AI). If the gains from AI continue to accrue disproportionately to a small group of entrepreneurs and investors while labor income declines, where will the demand for mass-produced goods come from? The central role of mass consumption in providing the scale of demand necessary to sustain investment and growth is encapsulated in a quote attributed to Henry Ford: "If you sell to the masses, you live with the classes; but if you sell to the classes, you live with the masses."

This paper studies the impact of income inequality on the composition of consumption demand and long-run economic growth. Our empirical analysis draws on a proprietary dataset from a U.S. fintech firm that combines income information with transaction-level consumption data for approximately 500,000 households. We match household expenditures to firms and combine these data with patent records to construct measures of the innovation intensity of the goods and services consumed by each household. Using these data, we show that the consumption bundle of the top 1 percent contains a markedly smaller share of goods produced by innovation-intensive firms than that of the remaining 99 percent of households.

One natural explanation of this pattern is that, as households enter the top of the income distribution, their consumption shifts toward luxury goods whose appeal derives less from functional attributes and more from what the marketing literature refers to as psychological value (Kotler (1967); Howard and Sheth (1969)). This value is often created through exclusivity and scarcity rather than through technological innovation. Goods such as the Hermès Birkin bag derive much of their value from limited availability, craftsmanship, and status signaling rather than superior functional performance.[2] As a result, consumption by the

[1] Walter Reuther, Selected Papers, January 1, 1961.

[2] Consistent with this view, Imas and Madarász (2024) show that consumers' willingness to pay increases

top 1 percent is less closely linked to mass-produced goods supplied by innovation-intensive firms.

This evidence suggests that reallocating income toward the top 1 percent reduces aggregate demand for innovation-intensive goods, dampens firms' incentives to innovate and expand output for mass consumption, and slows long-run economic growth.

We model the difference between the consumption of the top 1 percent and the remaining 99 percent by introducing a simple form of non-homothetic preferences that capture the idea of satiation. Utility from consuming each good is subject to a bliss point, beyond which additional units (another meal, another cell phone) provide no further utility. As households become richer, they therefore shift spending toward luxury, echoing the progression described by Maslow (1943) from satisfying basic needs to pursuing higher-level aspirations.

Growth remains possible even when households have reached the bliss point for existing goods. As in Romer (1990), innovators introduce new varieties that expand consumption possibilities and sustain growth.

The model yields two key results. When the bliss point is not binding, preferences are locally homothetic and changes in income inequality do not affect growth. When high-income households have reached the bliss point, however, reallocating income toward these households reduces demand for mass-market goods, lowers the profitability of creating new varieties, and slows economic growth.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 4 presents the empirical evidence. Section 5 develops the model and analyzes the potentially growth-reducing effects of rising income inequality associated with the AI revolution. Section 9 concludes.

2 Related Literature

Our paper relates to three strands of literature. First, it contributes to the research on inequality and growth, surveyed by Aghion et al. (1999). Existing theories emphasize channels operating through savings, capital accumulation, credit constraints, political economy, and innovation incentives. Beginning with Kuznets (1955), a large literature has studied the joint

when others are randomly excluded from access to the same good, providing direct evidence that exclusion itself can raise perceived value. Relatedly, Mandel and Imas (2026) show that human-generated art is valued partly because it is perceived as inherently more exclusive; even limited AI involvement weakens this perception and lowers valuations.

evolution of inequality and development, while more recent work has documented the rise in top income shares (Piketty and Saez, 2003) and explored its macroeconomic consequences. We complement this literature by focusing on the composition of consumption demand. We document that the consumption basket of the top 1 percent is less innovation-intensive than that of the remaining 99 percent and develop a mechanism through which this pattern links inequality to innovation and long-run growth.

Second, our paper relates to the literature on non-homothetic preferences, Engel curves, and structural transformation. Since Engel (1857), economists have documented systematic changes in consumption patterns as income rises (see e.g. Pollak, 1971 and Deaton and Muellbauer, 1980). A substantial body of empirical work shows that higher-income households consume different baskets of goods and tend to purchase higher-quality varieties (Verboven, 1996; Bils and Klenow, 2001; Hallak, 2006; Argente and Lee, 2021; Faber and Fally, 2022; Jaravel, 2019; Jaimovich et al., 2019). We contribute to this literature by documenting a new dimension of consumption heterogeneity: the innovation intensity of the goods and services consumed. We emphasize the distinction between innovation-intensive mass-market goods and exclusive goods whose value derives largely from scarcity and status.

This distinction echoes classical debates about luxury, though the eighteenth-century meaning of the term was broader than its modern association with exclusive status goods. David Hume viewed luxury, or "refinement in the arts," as potentially productive because it could stimulate employment (Hume, 1987a,b; see also Berry, 1994; Marshall, 2000). Adam Smith similarly recognized that the desire for wealth and distinction could create jobs but he was more ambivalent, warning that admiration for the rich rests partly on deception and can corrupt moral sentiments (Smith, 1982; Marshall, 2000). Our question is not whether high-income demand generates economic activity, but whether the activities it stimulates are innovation-intensive and scalable. In the modern economy, some forms of luxury demand are directed toward scarcity, status, and exclusivity rather than toward mass-market innovation. Our paper shows that this distinction matters for the link between inequality and long-run growth.

Our paper also relates to the literature on how non-homothetic preferences generate structural change (Kongsamut et al., 2001; Ngai and Pissarides, 2007; Acemoglu and Guerrieri, 2008; Duarte and Restuccia, 2010; Buera and Kaboski, 2012). We emphasize a distinct mechanism: once consumers become sufficiently affluent for the bliss point in ordinary consumption to bind, additional expenditure is increasingly directed toward luxury consump-

tion, weakening incentives to innovate in mass-market goods and thereby affecting long-run growth.

Third, our mechanism builds on a long tradition of models featuring satiation or bliss points, including Ramsey (1928), Samuelson and Solow (1956), Hall (1978), Mankiw (1982), Christiano et al. (1991), Heaton (1993), Cochrane (2014), Laitner (2000), Gollin et al. (2002), Gollin et al. (2007), and Foellmi and Zweimüller (2008). We show that growth can be sustained even when bliss points bind for all households. As consumers become satiated with existing goods, innovators create new varieties that expand consumption possibilities. This mechanism formalizes an idea implicit in the work of Kaldor (1966), Aglietta (1979), and Gordon (2016): long-run growth depends on the interaction between innovation and mass consumption. Our model highlights the limits of this interaction when income becomes concentrated among households whose consumption is less closely tied to innovation-intensive goods.

3 Data construction

Our analysis uses a new dataset compiled by a U.S. financial technology firm (henceforth fintech), merged with Compustat and patent data. The fintech specializes in financial data aggregation and analytics. It links securely, subject to user consent, accounts from thousands of financial institutions, including banks, credit cards, investment accounts, and loans, into a standardized format. This platform is widely used by fintech firms to construct consolidated measures of household financial activity and to inform lending decisions.

The construction of our dataset proceeds in five steps. First, we use transaction-level records to construct household measures of income and expenditures. Second, we identify the firms associated with each transaction and build household-specific expenditure portfolios over firms. Third, we match these firm identifiers to Compustat using a combination of machine-learning and manual validation procedures. Fourth, we merge firm-level patent data to the matched Compustat firms. Finally, we construct household-level measures of patent intensity by combining firm-level patent measures with household expenditure shares. The remainder of this section describes each step in detail.

Starting from raw transaction-level records containing several billion observations, we perform a sequence of cleaning, filtering, and matching steps to construct a household-income-expenditure-firm panel suitable for empirical analysis.

We first restrict expenditure measures to transactions classified as debits in clearly identifiable consumption categories, including restaurants, entertainment and recreation, groceries, electronics and general merchandise, automotive and fuel, utilities, telecommunications, education, rent, travel, healthcare, postage and shipping, gifts, pet care, home improvement, insurance, personal and family expenses, subscriptions, and ATM or cash withdrawals. This restriction ensures that measured outflows correspond closely to household consumption.

We then impose a set of sample selection criteria designed to ensure sufficient activity and reliable measurement. Households are required to have at least 40 transactions per month, including debits, credits, transfers, deposits, withdrawals, or card transactions, and to exhibit both positive total income and positive total consumption flows in each month included in the analysis. To further improve data quality, we remove duplicate transactions, restrict the sample to U.S. dollar-denominated accounts, trim extreme outliers, and exclude user identifiers associated with more than eight linked accounts, which are likely to correspond to small businesses rather than households.

After applying these cleaning, filtering, and matching procedures, we obtain a household-income-expenditure-firm panel that forms the basis of our empirical analysis. This dataset comprises 518,914 U.S. households and 2.1 billion individual transactions observed over the period 2015-2019 with an average transaction value of $65.

Households are expected to register all relevant financial accounts, allowing us to obtain a comprehensive view of household income and spending. Expenditures are constructed using outflows, including direct debits, one-off transfers, credit card spending, cash withdrawals, and mortgage payments. Income is constructed using inflows, including salaries and other regular income, tax refunds, interest income, and other investment and retirement income.

Household expenditures span approximately one million distinct firm identifiers, which allows us to construct household-specific firm expenditure portfolios. A subset of roughly 1,500 firms is publicly listed and can be matched to Compustat. These firms account for about 33 percent of all transactions in the data. For this subset, we measure patent intensityusing firm-level patent value. Restricting the analysis to Compustat-matched firms is unlikely to materially affect our results, as the bulk of innovation and patenting activity is concentrated among the large firms covered by Compustat.

In our sample, average annual household expenditure is approximately $51,172, which is about 14 percent lower than the corresponding average reported in the Consumer Expenditure Survey (CEX), $59,522. Average annual household income in our data is $82,492,

8.6 percent higher than the pre-tax income measure reported in the CEX, $75,870. These differences indicate the presence of sampling and measurement biases, but they are modest relative to what is typical in U.S. household-level datasets. Expenditures in our data account for approximately 86 percent of total CEX expenditures, compared with less than 20 percent coverage in the widely used Nielsen consumer panel.

Overall, we view the quality of the expenditure-income panel as high. The data have measurement error: some expenditures are missing, and income measures are noisy. But the coverage and internal consistency of the panel are substantially better than in most available U.S. datasets. Internal robustness checks suggest that the key fact we document is not materially affected by alternative filtering strategies.

We now discuss measurement issues related to the firm data. Our data are organized at the firm level rather than the product level, which may matter in the presence of multiproduct retailers such as Amazon and Walmart. In addition, roughly 25 percent of transactions in the raw data cannot be linked to a firm identifier. This unmatched share varies across the income distribution and is more pronounced at both the upper and lower tails. We examine the implications of this pattern and assess the robustness of our results to alternative treatments of unmatched transactions.

3.1 Matching Firms in the Fintech Dataset to Compustat

The fintech's transaction data include counterparty information that can be used to identify the firms receiving or paying funds, allowing us to construct an income–expenditure–firm panel. Two challenges arise. First, a nontrivial fraction of transactions lack counterparty information. Second, the fintech does not provide standardized firm identifiers: merchant names are free-text fields that are noisy, inconsistent, and often split into a primary and a secondary merchant name.

We first construct internal firm identifiers within the fintech. Using machine-learning text clustering methods, we group primary and secondary merchant names that refer to the same underlying firm. Specifically, we apply DBSCAN clustering to normalized merchant name strings, allowing for spelling variation and formatting differences. This procedure yields 1,032,164 distinct firm identifiers in the fintech's database.

To identify publicly listed firms, we build a bridge between the fintech's firm identifiers and Compustat. We use a transformer-based fuzzy name-matching algorithm (Arora and Dell (2023)'s "LinkTransformer") to compute semantic similarity scores between the mer-

chant names in the fintech database and Compustat firm names (`conm`). We retain only high-confidence matches, imposing a similarity threshold of 0.95.

Matching proceeds in stages. We first attempt to match the primary merchant names in the fintech database to Compustat firm names. When multiple Compustat firms match a given name in the fintech database, we manually review these cases and classify them as same entity, parent–subsidiary relationship, clearly different entities, product relationships, or ambiguous matches. We retain matches corresponding to the same entity or a parent–subsidiary relationship and drop all others. Remaining duplicates are resolved by selecting the highest-scoring match, yielding a one-to-one mapping.

We then repeat this procedure for secondary merchant names. The secondary name field is first split into multiple components using the delimiter "^", and each component is independently matched to Compustat using the same algorithm and filtering criteria.

Many transactions list payment processors or point-of-sale intermediaries rather than the underlying firm. We therefore manually identify and flag payment intermediaries (e.g., PayPal, Venmo, Square, Stripe, Klarna, Affirm, Apple Pay, Zelle). For each transaction, we assign the first non-intermediary firm match, prioritizing primary over secondary merchant names. Transactions for which all matched entities are intermediaries are retained but flagged accordingly.

Finally, we perform targeted manual validation to correct remaining false positives and false negatives, focusing on high-frequency merchants and categories prone to misclassification, notably car dealerships and common digital products (e.g., Audible transactions processed through Google Pay).

This procedure results in 1,684 distinct publicly listed firms in Compustat that can be reliably linked to the fintech's transaction data.

4 Empirical evidence

In this section, we provide new evidence on how the patent intensity of household consumption varies across the income distribution. We use patent intensity as an empirical proxy for innovation intensity. Patents are an imperfect measure of innovation: many innovations are not patented, patenting propensities vary across industries, and not all patents correspond to economically meaningful technological advances. Despite these limitations, patent-based measures provide a systematic and quantifiable indicator of innovative activity. Our use of

patent values from Kogan et al. (2017), which weight patents by their estimated economic importance, helps mitigate the concern that simple patent counts place excessive weight on low-value patents.

We begin by documenting patterns of income, expenditures, and saving across households. We then examine luxury consumption as a first indication of the relationship between income and innovation intensity before turning to measures based on households' complete consumption baskets.

4.1 Descriptive Evidence

For each household h, we construct total household income as the sum of all income inflows observed in the transaction data over the period 2015–2019. Income includes labor earnings and other recurring income sources recorded at the household level. To reduce transitory fluctuations and measurement error, we aggregate income over the full five-year window and work with

$$Y_h \equiv \sum_{t=2015}^{2019} \text{Income}_{h,t}.$$

Households are ranked according to $\log Y_h$, and distributional analyses are conducted either using income percentiles or bins of the income distribution.

Household expenditures are constructed analogously. For each household h, we compute total expenditures over the same five-year window,

$$E_h \equiv \sum_{t=2015}^{2019} \text{Expenditure}_{h,t},$$

excluding transactions associated with payment services, rent, and cash withdrawals. Aggregating expenditures over multiple years smooths large durable purchases and yields a stable measure of household consumption.

Table 1 shows the average income and expenditures across different income groups. The table documents a well-known empirical regularity: savings increase with income.

4.2 Luxury Consumption and Innovation

Our goal is to understand how the patent intensity of consumption varies across the expenditure distribution. As a first step, we focus on luxury goods. Luxury consumption is highly concentrated among households at the top of the expenditure distribution, making luxury

Table 1: Average Annual Income and Expenditures by Income Group

Income Group	Average Yearly Income	Average Yearly Expenditure
Overall	\$82,492	\$51,172
0–25%	\$28,068	\$42,557
25–50%	\$59,325	\$41,203
50–75%	\$85,619	\$48,356
75–90%	\$119,651	\$58,001
90–95%	\$156,550	\$69,837
95–99%	\$211,524	\$87,916
99–100%	\$500,225	\$243,425

firms a natural setting in which to examine whether the goods consumed by the rich are more or less innovation intensive than those consumed by the rest of the population.

Figure 1 plots the distribution of the expenditure share of the world's 100 largest luxury firms, as identified by Deloitte.[3] This share is approximately constant throughout most of the distribution and rises sharply among households near the very top. While the average household devotes only a small fraction of expenditure to these firms, the share increases rapidly among the highest-spending households. Luxury goods are disproportionately consumed by households at the top of the expenditure distribution.

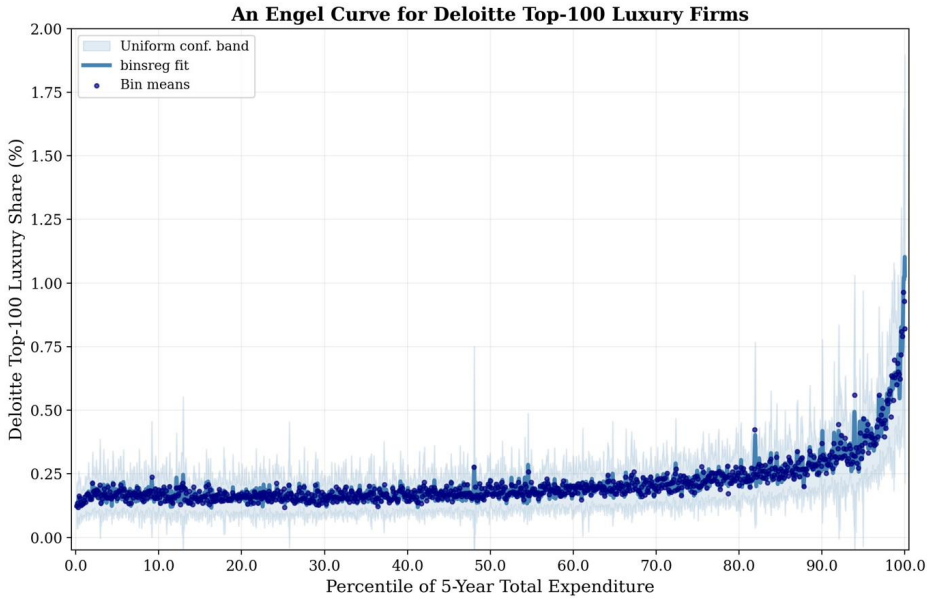


Figure 1: Expenditure share devoted to Deloitte's top-100 luxury firms.

[3]The classification is based on the sixth edition of Deloitte's *Global Powers of Luxury Goods* report, published in 2019.

We next match luxury firms to the universe of U.S. patents and compare their patent intensity with that of other firms. Figure 2 reports patent counts and citation-weighted patent counts. By both measures, luxury firms are substantially less innovation intensive than other firms. This evidence suggests that the goods consumed disproportionately by high-income households may embody less technological innovation than those consumed by middle-income households.

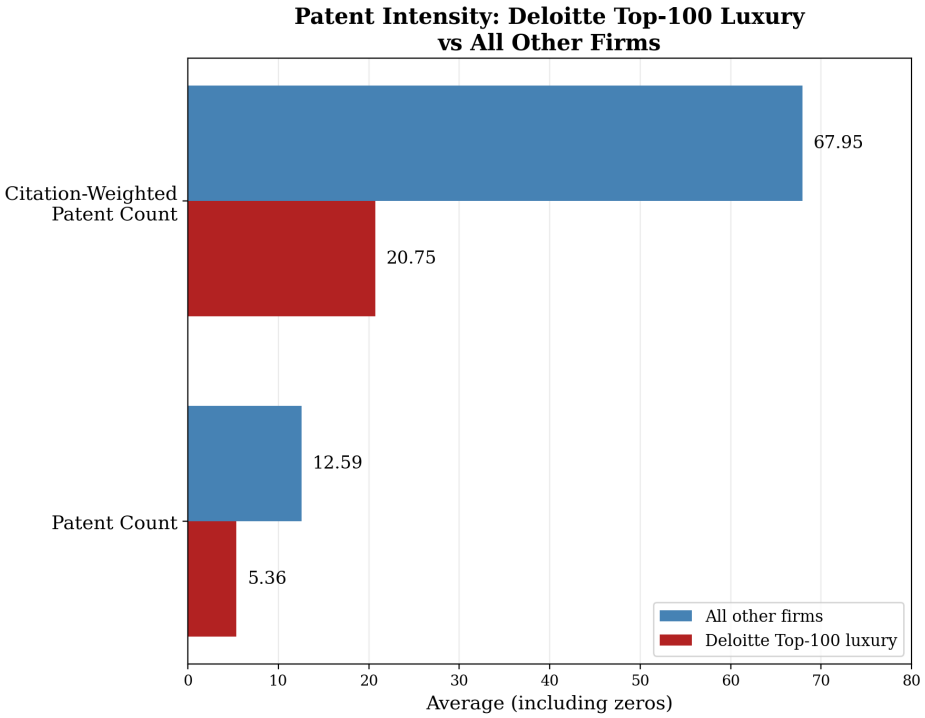


Figure 2: Patent intensity of Deloitte's top-100 luxury firms.

One concern is that the Deloitte list contains only a relatively small set of well-known luxury brands. To address this issue, we construct a broader, data-driven measure of luxury status. We classify a firm as luxury if more than one-half of its sales are made to households in the top 1 percent of the expenditure distribution. This procedure identifies 33,762 firms, representing 2.9 percent of all firms in our sample and accounting for 5.9 percent of total sales.

Figure 3 plots the expenditure share devoted to these firms. The resulting Engel curve is even more pronounced than the one based on the Deloitte classification. Households in the top 1 percent devote roughly 30 percent of their expenditure to firms classified as luxury, whereas the corresponding share for the remaining 99 percent of households is below 1 percent.

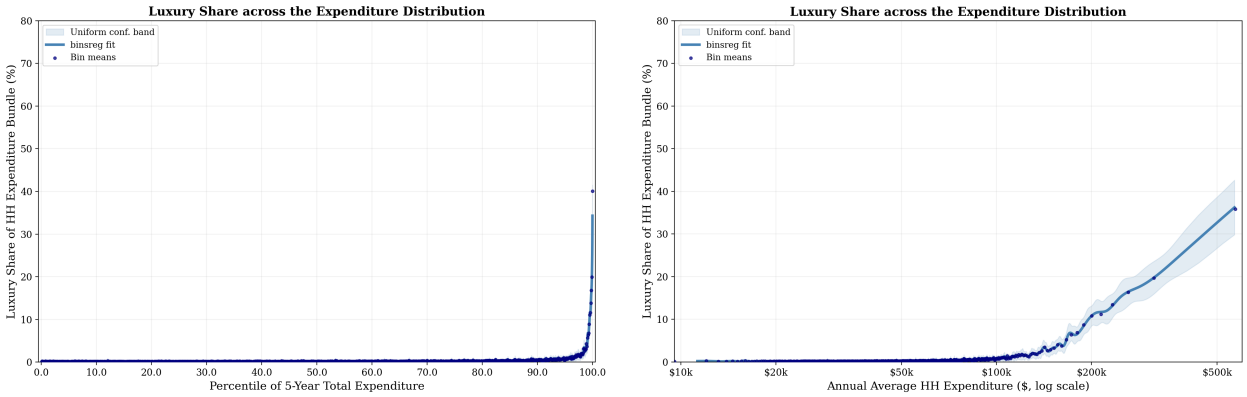


Figure 3: Expenditure share devoted to firms classified as luxury using the top-1-percent criterion.

Applying the same patent-matching procedure to this broader set of firms, we find that firms classified as luxury have fewer patents, fewer citation-weighted patents, and lower-valued patents than other firms (Figure 4). Thus, the negative association between luxury and innovation remains robust when we move from a curated list of luxury brands to a broader, data-driven classification.

4.3 Patent Intensity Across the Income Distribution

Taken together, our findings suggest that firms serving households at the very top of the expenditure distribution tend to be less innovative than firms serving the broader population. We now broaden the analysis, moving beyond luxury firms, and measure the patent intensity of the entire consumption basket across the expenditure distribution.

We first assign each firm a measure of innovation intensity based on its patent portfolio. We then combine these firm-level measures with household expenditure data to construct a household-specific measure of exposure to innovation. This procedure allows us to characterize the patent intensity of the goods and services consumed by households at different points of the income distribution.

Each expenditure transaction is associated with a firm identifier. Let e_{hf} denote total expenditures by household h on goods or services supplied by firm f over 2015–2019. The expenditure share of household h on firm f is given by

$$\omega_{hf} \equiv \frac{e_{hf}}{\sum_{f'} e_{hf'}}.$$

Patent data are merged at the firm level. To measure the innovation intensity of household consumption, we assign each firm f an innovation measure, M_f. We consider four alternative

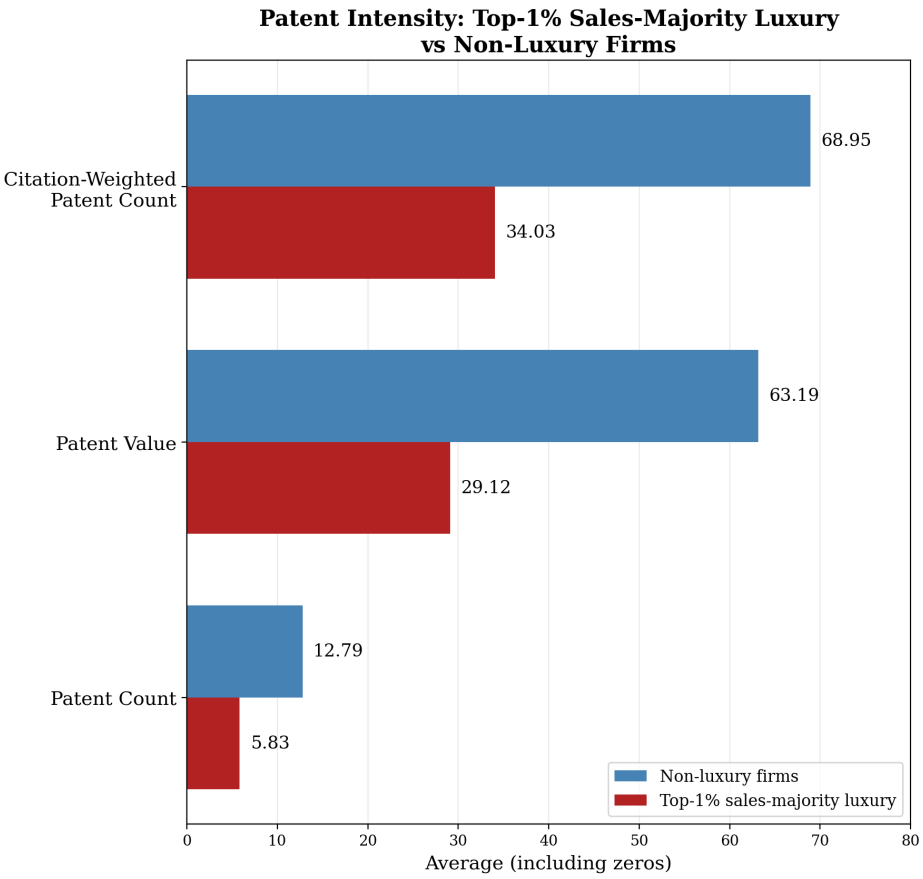


Figure 4: Patent intensity of firms classified as luxury.

measures: (i) patent counts, (ii) citation-weighted patent counts, (iii) patent values from Kogan, Papanikolaou, Seru, and Stoffman (2017), available for publicly-traded Compustat firms, and (iv) imputed patent values obtained by estimating the relationship between patent values and citation-weighted patent counts and using the resulting coefficients to extend patent value estimates to non-listed firms. These measures allow us to construct an Engel curve for innovation.

4.4 The Innovation Engel curve

We now examine how the innovation intensity of household consumption varies across the income distribution. We begin with the simplest measure of innovation—whether a firm holds any patents—and then consider progressively richer measures based on citation-weighted patent counts, patent values, and imputed patent values.

4.4.1 Patenting versus non-patenting firms

As a first pass, we compare expenditures on goods and services supplied by firms that hold at least one patent with expenditures on goods and services supplied by firms that hold no patents. Approximately 58 percent of the firms in our sample have no patents.

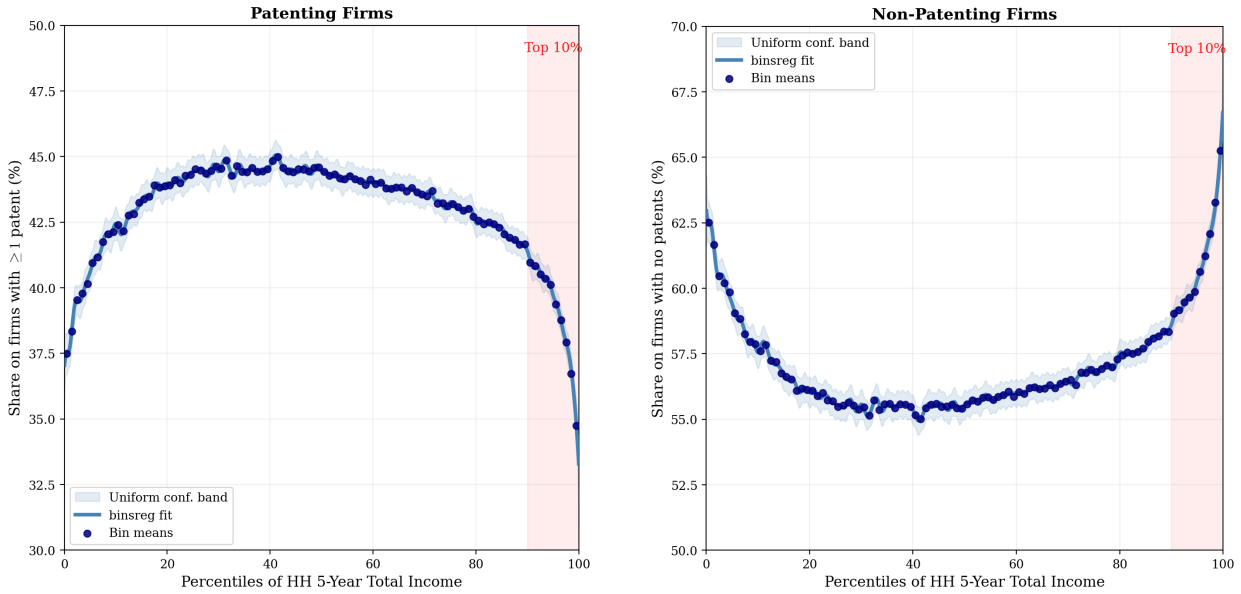


Figure 5: Income percentile (0–100) versus the share of consumption devoted to firms with at least one patent (left panel) and firms with no patents (right panel).

Figure 5 reveals a striking non-monotonic relationship between household income and the innovation intensity of consumption. The expenditure share directed toward patenting

firms initially rises with income but falls sharply at the top, while the share directed toward firms without patents displays the opposite pattern.

For the top income decile, the share of spending on patenting firms declines from approximately 45 to 35 percent, with a corresponding increase from 55 to 65 percent for firms without patents. Thus, the consumption bundles of the highest-income households are less innovation-intensive than those of middle-income households.

4.4.2 Citation-weighted patent counts

The distinction between patenting and non-patenting firms is coarse because patents differ greatly in their economic importance. We therefore construct a more refined measure of innovation intensity based on citation-weighted patent counts. For household h, this measure is

$$\mathrm{CW}_h \equiv \sum_f \omega_{hf} \, \mathrm{CW}_f,$$

where ω_{hf} denotes household h's expenditure share at firm f and CW_f is the firm's citation-weighted patent count.

Figure 6 plots this measure across the income distribution. The innovation intensity of household consumption declines sharply at the top: relative to households in the bottom 99 percent, those in the top 1 percent consume expenditure bundles that are approximately 35 percent less innovation-intensive, as measured by citation-weighted patent counts. This result confirms that the decline in patent intensity documented in Figure 5 is not driven merely by differences in patent counts; it persists when patents are weighted by their influence.

4.4.3 Patent value

Citation-weighted patent counts remain an imperfect proxy for innovation because they do not directly measure economic value. We therefore turn to the patent values estimated by Kogan, Papanikolaou, Seru, and Stoffman (2017). These values are inferred from the issuing firm's stock-market reaction around patent grant dates, adjusting for movements in stock prices unrelated to the patent announcement.

For each household h, we compute the patent value of its expenditure bundle as

$$\mathrm{PV}_h \equiv \sum_f \omega_{hf} \, \mathrm{PV}_f,$$

where PV_f denotes the patent value of firm f.

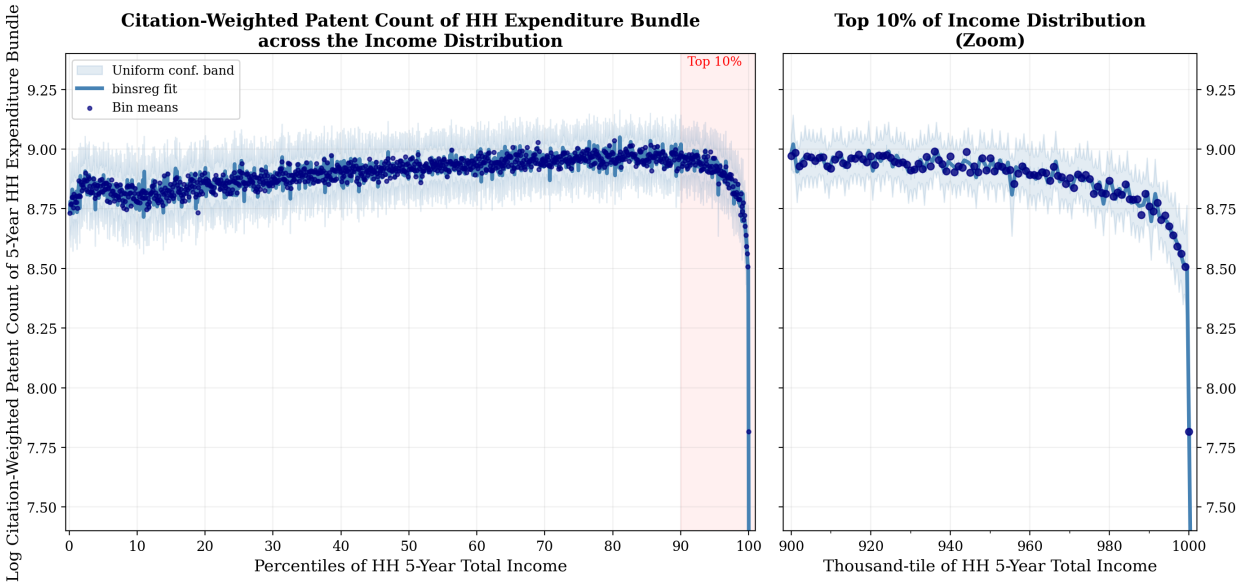


Figure 6: Income percentile (0–100) versus expenditure innovation intensity measured by citation-weighted patent counts.

Figure 7 plots this measure across the income distribution. Once again, innovation intensity declines at the top. Households in the top 1 percent consume expenditure bundles that are approximately 22 percent less innovation-intensive than those of the remaining 99 percent.

4.4.4 Imputed patent value

The drawback of the Kogan et al. (2017) measure of patent value is that it is only available for publicly traded firms. To extend the analysis beyond this set of firms, we construct an imputed patent-value measure. Using Compustat data for publicly traded firms, we estimate the relationship between the logarithm of patent value and a quadratic function of the logarithm of citation-weighted patent counts. We then use this relationship to predict patent values for non-Compustat firms, under the assumption that citation-weighted patents map into economic patent value in a similar way for listed and non-listed firms.

Figure 8 presents our results using the resulting measure. The pattern is again similar: households at the top of the income distribution consume bundles with 30 percent lower innovation intensity than those of other households. Relative to the bottom 99 percent, households in the top 1 percent consume expenditure bundles that are approximately 30 percent less innovation-intensive when innovation is measured using imputed patent values.

The results are remarkably consistent across all measures of innovation intensity. Firms

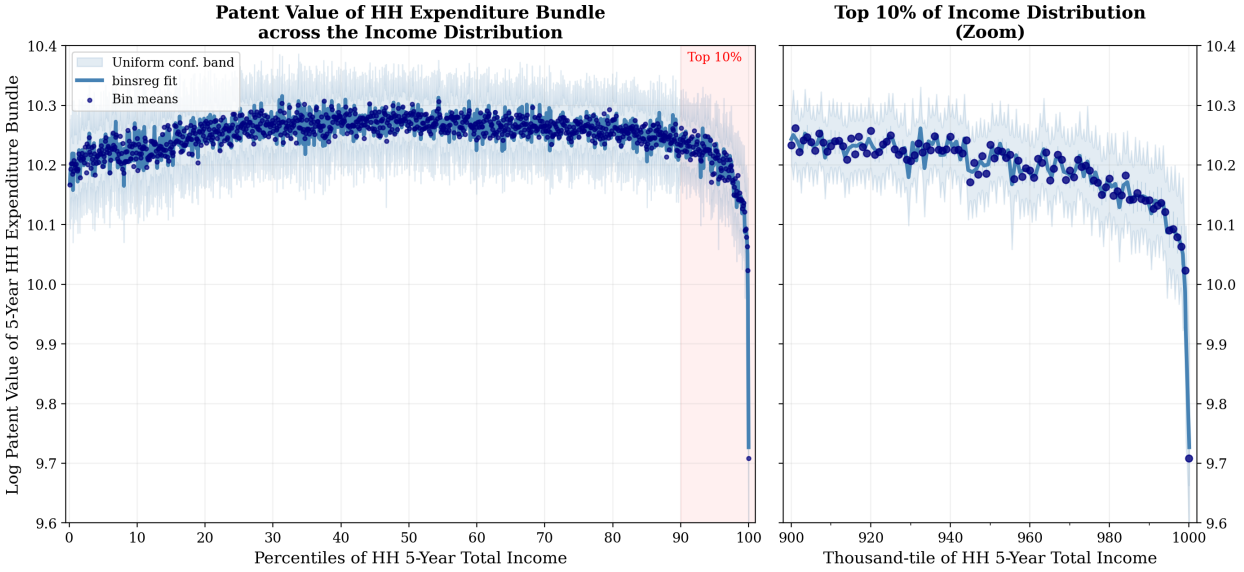


Figure 7: Income percentile (0–100) versus expenditure innovation intensity measured by patent value.

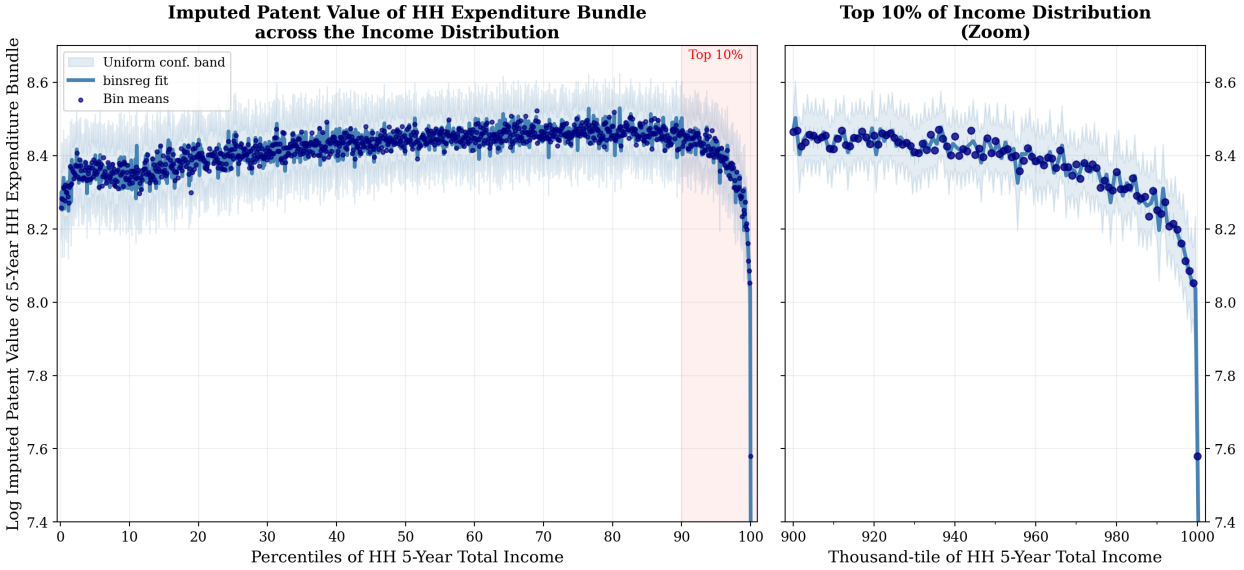


Figure 8: Income percentile (0–100) versus expenditure innovation intensity measured by imputed patent value.

that disproportionately sell to the highest-income households are systematically less innovation-intensive. Relative to the bottom 99 percent, households in the top 1 percent devote a larger share of their expenditure to firms with no patents, firms with lower citation-weighted patent counts, and firms with lower patent values. Consequently, rising inequality shifts demand away from innovation-intensive goods and services.

Schumpeter (1950) and Hayek (1960) emphasized the role of high-income households as early adopters of new technologies. Their willingness to purchase new goods at high initial prices may support experimentation and early diffusion. But early adoption does not generate the broad demand on which sustained profitability ultimately rests. Our results suggest that demand by the bottom 99 percent of the income distribution is essential for creating the scale necessary to make innovation profitable.

5 Model

We study how rising income inequality affects innovation and growth when high-income households have already satisfied their demand for existing mass-produced goods. Additional income accruing to these households is therefore directed toward luxury consumption, reshaping the demand for mass-produced goods and the allocation of productive resources. We trace how these shifts affect prices, profits, and the incentives to create new goods.

We characterize balanced growth paths (BGPs), on which quantities and prices grow at constant rates, and omit time arguments when doing so does not reduce clarity.

The structure of the model is as follows. There are two types of consumption goods: differentiated mass-produced goods and premium, or luxury, goods. Each mass-produced variety has a bliss point, $\bar{c}$, beyond which additional consumption yields no further utility. There is nevertheless no satiation in overall consumption: households value new varieties of mass-produced goods as well as additional luxury consumption. As we discuss in Section 5.1.4, this simple departure from homothetic preferences generates Engel curves consistent with our empirical work.

Competitive firms produce materials using labor and differentiated intermediate inputs. Monopolists produce these intermediates using materials and sell them to both materials producers and households. Growth arises from the introduction of new intermediate varieties, as in Romer (1990). New varieties expand consumption opportunities even when households have reached bliss on existing goods.

5.1 Households

There is a unit measure of households. A fraction $m_H \in (0,1)$ have productivity ϵ_H, and the remaining fraction $m_L = 1 - m_H$ have productivity ϵ_L, with $0 < \epsilon_L < \epsilon_H$. Labor supply is inelastic, and aggregate efficiency labor is normalized to one:

$$m_H \epsilon_H + m_L \epsilon_L = 1. \tag{1}$$

Let $s \in (m_H, 1)$ denote the high type's share of aggregate efficiency labor. The normalization above implies

$$\epsilon_H = \frac{s}{m_H}. \tag{2}$$

$$\epsilon_L = \frac{1-s}{m_L}. \tag{3}$$

Holding the population shares m_H and m_L fixed, an increase in s constitutes a mean-preserving spread of the productivity distribution: it raises ϵ_H and lowers ϵ_L while leaving aggregate efficiency labor unchanged. Equilibrium wages, profits, aggregate expenditure, and growth can respond to this redistribution.

5.1.1 Preferences and the bliss point

Both types have the same preferences, with $0 < \gamma < 1$ and $\rho > 0$. A household of type $j \in \{L, H\}$ has lifetime utility

$$\int_0^\infty e^{-\rho t} \left[\gamma \log C_j(t) + (1-\gamma) \log q_j(t)\right] dt. \tag{4}$$

where q_j is the household's luxury consumption and C_j is an aggregate of mass-produced varieties:

$$C_j = \left[\int_0^n \min\{c_{i,j}, \bar{c}\}^{1/\nu} \, di\right]^\nu. \tag{5}$$

Here $n(t)$ is the measure of available varieties, $c_{i,j}$ is type-j consumption of variety i, and $\bar{c} > 0$ is the common bliss point. The parameter γ is the utility weight on mass-produced consumption, ρ is the rate of time preference, and $\nu > 1$ governs substitution across varieties. Once a household consumes $\bar{c}$ units of a variety, additional units provide no utility. Below bliss, the elasticity of substitution across varieties is

$$\sigma = \frac{\nu}{\nu - 1}. \tag{6}$$

5.1.2 The Household Budget Constraint

Let w denote the wage per efficiency unit of labor, p_i the price of mass-produced variety i, and P_Q the price of luxury consumption. A type-j household earns labor income $w\epsilon_j$ and chooses consumption and asset holdings subject to

$$\begin{aligned}\int_0^n p_{s,i}\dot{s}_{i,j}\,di + p_s^q\dot{s}_{q,j} + \dot{b}_j + \mathcal{I}_j^N = &\int_0^n s_{i,j}\pi_i\,di + s_{q,j}\Pi_q + rb_j + w\epsilon_j \\ &- \int_0^n p_i c_{i,j}\,di - P_Q q_j.\end{aligned} \tag{7}$$

The variable b_j denotes holdings of a risk-free bond paying interest rate r. Households may also hold equity claims in two types of firms. The variable $s_i(\epsilon)$ denotes the holdings of the household ϵ in the intermediate-goods producer i, which generates profits π_i. The variable $s_q(\epsilon)$ denotes holdings of shares in luxury producers, which generate total profits π_q and trade at price p_s^q. The variables $\dot{s}_{i,j}$ and $\dot{s}_{q,j}$ represent transactions in existing claims, while $\mathcal{I}_j^N$ denotes purchases of newly issued equity by intermediate goods producers. Asset choices satisfy the usual no-Ponzi and transversality conditions. To simplify, we assume that shares in intermediate-goods producers and luxury firms are distributed across households in proportion to ϵ.

Total consumption expenditure is

$$\mathcal{E}_j = \int_0^n p_i c_{i,j}\,di + P_Q q_j. \tag{8}$$

Initial equity ownership is proportional to a household's efficiency units of labor, and households initially hold no bonds.[4] This initial distribution supports the balanced growth allocation studied below. Section 6.1 derives the associated consumption expenditure, accounting for the saving that finances new varieties.

5.1.3 Consumption Allocation

When consumption lies below the bliss point, demand has a constant elasticity. Once a household reaches the bliss point, however, its demand becomes locally inelastic. The following demand system captures both cases. Let λ_j denote the current-value marginal utility of wealth for household type j. The CES price index is given by

$$P_C = \left[\int_0^n p_i^{1-\sigma}\,di\right]^{1/(1-\sigma)}. \tag{9}$$

[4] Linking ownership to productivity allows a single parameter to govern the distribution of both labor and capital income.

For each household type $j \in \{L, H\}$, optimal demand has a Dixit–Stiglitz form, capped at the bliss point:

$$c_{i,j} = \min\left\{\left(\frac{p_i}{P_C}\right)^{-\sigma} C_j^d, \bar{c}\right\}, \qquad i \in [0, n], \tag{10}$$

The demand scale C_j^d determines the level of purchases before imposing the cap. It is determined jointly with actual composite consumption:

$$C_j^d = \left(\frac{\gamma}{\lambda_j P_C}\right)^{\sigma} C_j^{1-\sigma}. \tag{11}$$

Luxury consumption satisfies

$$P_Q q_j = \frac{1-\gamma}{\lambda_j}. \tag{12}$$

When consumption is below the bliss point for every variety, $C_j^d = C_j$ and expenditure on mass-produced goods satisfies

$$P_C C_j = \gamma \mathcal{E}_j. \tag{13}$$

When the bliss point is binding, C_j^d generally differs from actual composite consumption C_j. Expenditure on luxury is the remainder of the consumption budget:

$$P_Q q_j = \mathcal{E}_j - \int_0^n p_i c_{i,j}\, di. \tag{14}$$

Aggregate household demand is obtained by weighting each household's purchases by its population share. We evaluate these demands after determining expenditure on the BGP in Section 6.2.

Finally, consumption expenditure is determined jointly with saving. Households equate the marginal value of spending today with the value of postponing that expenditure. The co-state equation associated with bond holdings is given by

$$\frac{\dot{\lambda}_j}{\lambda_j} = \rho - r. \tag{15}$$

5.1.4 Engel Curves

In an equilibrium with symmetric prices, $p_i = P$, let c_j denote type-j consumption of each variety. A household reaches the bliss point for every variety when its expenditure on mass-produced goods equals or exceeds

$$\mathcal{E}^* = \frac{nP\bar{c}}{\gamma}.$$

For either type $j \in \{L, H\}$, expenditure below this threshold is allocated in constant shares: γ to mass-produced goods and $1-\gamma$ to luxuries. Expenditure on both categories therefore rises proportionally with total spending.

Above $\mathcal{E}^*$, consumption of each mass-produced variety equals $\bar{c}$. At given prices and a given number of varieties, expenditure on these goods remains fixed at $nP\bar{c}$, and all additional expenditure goes to luxuries. The luxury expenditure share is therefore

$$\frac{P_Q q_j}{\mathcal{E}_j} = \max\left\{1-\gamma,\, 1-\frac{nP\bar{c}}{\mathcal{E}_j}\right\}. \tag{16}$$

This expenditure share is constant below $\mathcal{E}^*$ and rises once expenditure exceeds that threshold, producing a kinked Engel curve like the one shown in Figure 3. Thus, the bliss point generates the observed variation in expenditure shares even though households have identical preferences.

These differences in marginal spending patterns are the source of the redistribution effect. Holding prices and aggregate expenditure fixed, reallocating one dollar from the low-income type to the high-income type shifts γ dollars of expenditure from mass-produced goods to luxuries. The equilibrium analysis below allows both prices and aggregate expenditure to adjust endogenously.

5.2 Production and Pricing

Intermediate producers sell directly to households and also supply inputs to materials firms. Materials, produced using labor and intermediate goods, are used both to manufacture existing goods and to create new varieties. Because luxury services also require labor, their expansion draws workers away from materials production.

Materials. Materials serve as the numeraire, with price $P_M = 1$. Let M denote aggregate materials output and x_i the materials sector's input of variety i. Competitive firms combine these inputs with efficiency labor L_M according to

$$M = L_M^{\alpha} \int_0^n x_i^{1-\alpha}\, di. \tag{17}$$

For $\alpha \in (0,1)$, profit maximization gives intermediate demand

$$x_i = L_M \left(\frac{1-\alpha}{p_i}\right)^{1/\alpha}. \tag{18}$$

At symmetric prices, let x_m denote materials producers' demand for each variety. Then

$$x_m = L_M \left(\frac{1-\alpha}{P} \right)^{1/\alpha}. \tag{19}$$

The wage is

$$w = n\alpha \left(\frac{1-\alpha}{P} \right)^{(1-\alpha)/\alpha}. \tag{20}$$

For tractability, we impose

$$\alpha = \frac{\nu - 1}{\nu}. \tag{21}$$

Demand from materials producers and from households below the bliss point therefore has the same price elasticity, σ. The share of sales to satiated households is the only source of variation in the elasticity relevant for intermediate-goods pricing.

Luxury services. Competitive assemblers combine a unit measure of differentiated services, indexed by ℓ, to produce luxury output Q. Let $x_{q,\ell}$ denote the quantity of service ℓ used in assembly. The production function is

$$Q = \left[\int_0^1 x_{q,\ell}^{1-\beta} \, d\ell \right]^{1/(1-\beta)}. \tag{22}$$

The parameter $\beta \in (0,1)$ governs substitution across services. The price $p_{q,\ell}$ of service ℓ satisfies the assembler's inverse demand condition:

$$p_{q,\ell} = P_Q \left(\frac{Q}{x_{q,\ell}} \right)^{\beta}. \tag{23}$$

Each service is produced by a monopolist using efficiency labor one-for-one. Its demand elasticity is $1/\beta$, so the gross markup over the wage is $1/(1-\beta)$. Under symmetry, each service producer supplies Q, and the sector uses Q efficiency units of labor. The price of the luxury good is

$$P_Q = \frac{w}{1-\beta}. \tag{24}$$

Competitive assemblers earn zero profits. The total profits of service producers are:

$$\Pi_q = \frac{\beta}{1-\beta} wQ. \tag{25}$$

These profits, which are distributed to households as dividends, influence equilibrium consumption expenditure.

Intermediate goods. A monopolist produces each intermediate variety using $\kappa > 0$ units of materials per unit of output. We assume that price discrimination is not possible, so the monopolist must charge the same price to households and materials producers. Consider the regime in which the low-income type is below the bliss point and the high-income type has reached it. Let C^u denote aggregate low-income type consumption of each variety:

$$C^u = m_L c_L, \tag{26}$$

Aggregate high-income type consumption per variety is

$$B = m_H \bar{c}, \tag{27}$$

The price-sensitive component of demand for intermediate goods includes purchases by materials firms and consumption by low-income households:

$$\mathcal{R} = x_m + C^u. \tag{28}$$

Total sales of an intermediate variety, denoted by D, are the sum of this elastic component, $\mathcal{R}$, and the inelastic demand of satiated households, B:

Total sales, D, of an intermediate variety is the sum of the elastic demand, $\mathcal{R}$, and the inelastic demand, B,

$$D = \mathcal{R} + B. \tag{29}$$

The profit of one intermediate producer is

$$\pi_x = (P - \kappa)D. \tag{30}$$

As shown in the Appendix, a necessary first-order condition for an intermediate-goods producer's problem in a symmetric equilibrium is

$$P = \frac{\nu\kappa}{1 - (\nu - 1)B/\mathcal{R}}. \tag{31}$$

When the bliss point binds for neither household type, the inelastic component of demand, B, is zero, and demand has elasticity σ. The producer therefore charges the standard constant-markup price, $\nu\kappa$. When the high-income type has reached the bliss point, by contrast, its demand contributes to sales but is locally price inelastic, raising the gross markup above ν.[5]

[5]Every smooth monopoly optimum in this regime also satisfies

$$\nu\kappa < P < (2\nu - 1)\kappa. \tag{32}$$

We discuss in the Appendix how this bound determines the effect of prices on innovation profits.

5.3 New Varieties

New varieties are created in the intermediate-goods sector, with each variety generating a new stream of profits. Creating a variety requires $1/\Psi$ units of materials, where $\Psi > 0$ measures research productivity. Let m denote the quantity of materials devoted to research and development. The number of varieties then evolves according to

$$\dot{n} = \Psi m. \tag{33}$$

A new variety grants its owner a perpetual patent to produce that good. Let V denote the value of an intermediate-goods patent. Because each firm has a single equity claim outstanding, V is also the firm's share price in a symmetric equilibrium. When innovation is strictly positive, free entry equates the value of a patent to the cost of creating a new variety:

$$V = \frac{1}{\Psi}. \tag{34}$$

No arbitrage requires

$$rV = \pi_x + \dot{V}. \tag{35}$$

Since free entry fixes the patent value in units of materials, positive innovation implies

$$r = \Psi \pi_x. \tag{36}$$

6 Equilibrium

We determine household expenditure, sectoral production, and pricing jointly. We first describe the balanced growth allocation and its expenditure accounting, then impose market clearing and characterize the equilibrium price.

6.1 Balanced Growth and Household Expenditure

Growth and normalized prices. We study symmetric BGPs with positive innovation and an initial measure of varieties $n(0) > 0$. The rate of variety growth is

$$g = \frac{\dot{n}}{n}. \tag{37}$$

Along a BGP, P, consumption per variety c_j, and luxury quantities q_j are constant. Wages, the luxury price, luxury dividends, and household consumption expenditure grow with n.

We divide these growing variables by the number of varieties to obtain stationary quantities. In particular, the normalized wage is

$$W \equiv \frac{w}{n}. \tag{38}$$

The normalized luxury price is

$$p_q \equiv \frac{P_Q}{n}. \tag{39}$$

Normalized aggregate luxury profits are

$$\pi_q \equiv \frac{\Pi_q}{n}. \tag{40}$$

Equations (20), (24), and (25) imply

$$W(P) = \alpha \left(\frac{1-\alpha}{P} \right)^{(1-\alpha)/\alpha}. \tag{41}$$

$$p_q = (1+\delta) W(P). \tag{42}$$

$$\pi_q = \delta W(P) Q. \tag{43}$$

The distinction between consumption per-variety and aggregate consumption matters for growth. In a symmetric equilibrium, $C_j = n^{\nu} c_j$ and $P_C = n^{1-\nu} P$. Thus, a household can remain at the bliss point for each existing variety while its consumption aggregate grows at rate νg as new varieties become available.

Since q_j is constant and P_Q grows at rate g, the luxury first-order condition (12) implies that λ_j falls at rate g. The co-state condition (15) therefore yields

$$r = \rho + g. \tag{44}$$

Combining this condition with free entry and intermediate profits yields

$$g = \Psi(P - \kappa) D - \rho. \tag{45}$$

Profits from one variety therefore determine the return to innovation and the rate at which the product range expands. We discuss the formal proof of this result in in Appendix.

6.2 Consumption by Household Type

To derive consumption by the two household types, we first determine the consumption budget that remains after financing innovation. This budget combines labor income, net dividends from intermediate-goods firms, and dividends from luxury producers. Income

generated in the luxury sector therefore feeds back into the demand for mass-produced goods.

Let Z denote aggregate household consumption expenditure per existing variety:

$$Z \equiv \frac{m_L \mathcal{E}_L + m_H \mathcal{E}_H}{n}. \tag{46}$$

The numerator includes spending on both mass-produced goods and luxuries. Consequently, nZ is total consumption expenditure in the economy. Savings, which in equilibrium finance investment in new varieties, are excluded from this consumption measure. As we show in the Appendix, household expenditure satisfies

$$\mathcal{E}_j = nZ\epsilon_j, \tag{47}$$

where

$$Z = W + \frac{\rho}{\Psi} + \pi_q. \tag{48}$$

We can now express the purchases of the two household types in terms of the expenditure levels derived above. Our main analysis focuses on the mixed regime, in which expenditure by the low-income type lies below the threshold $\mathcal{E}^*$ defined in Section 5.1.4, while expenditure by the high-income type lies strictly above it. Using (47), this regime is characterized by

$$\frac{\gamma Z \epsilon_L}{P} < \bar{c} < \frac{\gamma Z \epsilon_H}{P}. \tag{49}$$

These inequalities determine the position of each household type relative to the common bliss point. The low-income type's consumption of each variety is

$$c_L = \frac{\gamma Z \epsilon_L}{P}. \tag{50}$$

Its normalized luxury expenditure is

$$p_q q_L = (1-\gamma) Z \epsilon_L. \tag{51}$$

The high-income type consumes each variety at the bliss point:

$$c_H = \bar{c}. \tag{52}$$

It allocates the remainder of its consumption budget to luxury goods:

$$p_q q_H = Z\epsilon_H - P\bar{c}. \tag{53}$$

Weighting the low-income type's consumption by its population share yields the price-sensitive component of household demand defined above:

$$C^u = \frac{\gamma(1-s)Z}{P}. \tag{54}$$

Aggregate household demand per variety is therefore

$$m_L c_L + m_H c_H = \frac{\gamma(1-s)Z}{P} + B. \tag{55}$$

The first term varies with expenditure by the low-income type, whereas B remains fixed as long as the high-income type remains at the bliss point. This asymmetry is the source of the income redistribution effect on the demand for intermediate goods.

6.3 Market Clearing

Labor. Efficiency units of labor are employed in materials production and luxury services. Because the aggregate supply of efficiency labor is normalized to one, labor-market clearing requires

$$L_M + Q = 1. \tag{56}$$

Materials are used both in intermediate-goods production and in research. Substituting these uses into (19) yields materials producers' demand for each intermediate variety:

$$x_m = (1-Q)\left(\frac{1-\alpha}{P}\right)^{1/\alpha}. \tag{57}$$

Luxury. Market clearing in the luxury sector requires output to equal aggregate household consumption:

$$Q = m_L q_L + m_H q_H. \tag{58}$$

Multiplying by the normalized price of luxury services and substituting the expenditure equations for the two household types yields

$$p_q Q = Z(1-\gamma+\gamma s) - PB. \tag{59}$$

This condition links the composition of household expenditure to employment in the luxury sector, since each unit of luxury output requires one efficiency unit of labor.

Intermediate goods. The quantity D sold by each intermediate producer is the sum of demand from materials firms and the two components of household demand:

$$D = x_m + C^u + B. \tag{60}$$

Materials. Let $Y_M = M/n$ denote materials output per existing variety. Under symmetry,

$$Y_M = L_M^{\alpha} x_m^{1-\alpha}. \tag{61}$$

Competition in the materials sector implies the zero-profit condition $Y_M = WL_M + Px_m$. Materials are used as inputs in both intermediate-goods production and the creation of new varieties. Because research requires $m/n = g/\Psi$ units of materials per existing variety, market clearing for materials requires

$$Y_M = \kappa D + \frac{g}{\Psi}. \tag{62}$$

The two terms on the right-hand side represent materials used in current production and investment in new varieties.

Bonds are in zero net supply, while the proportional allocation of equity holdings across housholds ensures that the equity market clears.

6.4 Equilibrium Characterization

Definition 1 (Interior equilibrium) *A symmetric interior balanced growth path (BGP) is an equilibrium in which $n(t) = n(0)e^{gt}$, households and firms optimize, asset prices satisfy the no-arbitrage and free-entry conditions, and all markets clear. Along this path, normalized allocations and prices are constant, $g > 0$, $0 < Q < 1$, and the mixed-regime condition* (49) *holds: low-income households remain below the bliss point, while high-income households have reached it. Each intermediate producer chooses a price that globally maximizes its profits.*

The equilibrium can be characterized by comparing two quantities. For a given intermediate-goods price and income distribution, market clearing and household expenditure determine the elastic component of intermediate demand. The monopolist's pricing condition, in turn, determines the level of elastic demand consistent with that price. In equilibrium, these two quantities coincide, subject to the allocation and optimality conditions stated below.

First, combining luxury-market clearing with (48) and the equations for the price and profits of luxury services yields

$$Q(P, s) = \frac{[1 - \gamma + \gamma s][\rho/\Psi + W(P)] - PB}{W(P)[1 + \delta\gamma(1 - s)]}. \tag{63}$$

Substituting this expression for output into the expenditure identity yields the normalized aggregate consumption expenditure defined in (46):

$$Z(P, s) = \frac{\rho}{\Psi} + W(P)[1 + \delta Q(P, s)]. \tag{64}$$

The dependence of Z on P and s reflects how wage income and luxury-sector dividends respond to prices and redistribution. Adding materials demand to low-income household consumption then yields

$$\mathcal{R}(P,s) = [1 - Q(P,s)] \left(\frac{1-\alpha}{P}\right)^{1/\alpha} + \frac{\gamma(1-s)Z(P,s)}{P}. \tag{65}$$

The pricing condition implies that elastic demand satisfies

$$\mathcal{R}^{\text{req}}(P) = \frac{(\nu-1)BP}{P - \nu\kappa}. \tag{66}$$

Hence, the equilibrium price solves

$$F(P,s) \equiv \mathcal{R}(P,s) - \mathcal{R}^{\text{req}}(P), \tag{67}$$

$$F(P,s) = 0. \tag{68}$$

This condition reflects the effects of wages, luxury production, and household income on the equilibrium price. The Appendix provides a formal derivation of this condition.

7 Inequality and Growth

We vary s while holding population shares and all other primitives fixed. We begin with the benchmark in which both household types remain below the bliss point and redistribution does not affect growth. We then turn to the mixed regime, decomposing the growth response into two components: the effect of prices on profits and innovation, and the effect of redistribution on the equilibrium price.

7.1 The Benchmark with Unconstrained Households

When both household types allocate the same shares of their resources to each consumption category, redistributing expenditure between them leaves aggregate demand unchanged. The equilibrium analysis below shows that this conclusion continues to hold when wages and profits are determined endogenously.

Proposition 2 (Inequality neutrality) *Along a BGP with positive innovation, any change in s under which both household types remain strictly below the bliss point leaves aggregate allocations, prices, and growth unchanged.*

Aggregate household demand for each variety is then $\gamma Z/P$, and every component of demand has price elasticity σ. The equilibrium price is therefore

$$P_0 = \nu\kappa. \tag{69}$$

Let $W_0 = W(P_0)$. Luxury output satisfies

$$Q_0 = \frac{(1-\gamma)(\rho/\Psi + W_0)}{W_0(1+\delta\gamma)}. \tag{70}$$

Normalized expenditure is

$$Z_0 = \frac{(1+\delta)(\rho/\Psi + W_0)}{1+\delta\gamma}. \tag{71}$$

Neither expression depends on s. Intermediate-goods demand, profits, and growth are therefore invariant to redistribution. The mixed regime departs from this benchmark because high-income households direct all marginal expenditure toward luxury services.

7.2 The Effect of Inequality on Growth

Our object of interest is dg/ds, the effect of redistributing income toward the high-income type on the rate at which new varieties are created. We focus on the mixed regime, in which low-income households consume below the bliss point while high-income households consume each variety at bliss. Holding population shares and the bliss point fixed, aggregate high-income consumption per variety, B, remains unchanged. The pricing condition then allows profits, and hence growth, to be expressed as a function of the equilibrium price of intermediate goods. We can therefore write

$$\frac{dg}{ds} = \frac{dg}{dP}\frac{dP}{ds}. \tag{72}$$

We proceed in two steps. First, we examine how innovation and the price of intermediate goods move together, as summarized by dg/dP. Second, we determine how redistribution affects the equilibrium price, dP/ds and then comnine these two terms to obtain the effect of inequality on growth.

7.2.1 Prices, Profits, and Growth

Consider first dg/dP. The value of a new variety depends on the profits it generates, which are determined by both the producer's margin and its sales. The pricing condition links these two components: a higher optimal price requires elastic demand to be smaller relative

to the fixed quantity purchased by households at the bliss point. Substituting this condition into (45) yields

$$g(P) = \Psi\nu B\frac{(P-\kappa)^2}{P-\nu\kappa} - \rho. \tag{73}$$

A higher price has two opposing effects on profits. It increases the margin on each unit sold but reduces sales to price-sensitive buyers, including both low-income households and materials producers.

The monopolist's optimality condition implies that the decline in sales dominates the increase in the profit margin. The higher markup is an optimal response to weaker price-sensitive demand: repricing mitigates the resulting loss of profits but cannot eliminate it. Consequently,

$$\frac{dg}{dP} < 0. \tag{74}$$

This result compares allocations along the locus defined by the pricing condition. In the Appendix we derive the corresponding expressions for sales and profits and establishes the sign using the monopolist's optimality bound.

7.2.2 Redistribution and Equilibrium Prices

We next turn to dP/ds. To isolate the shift in demand, hold the price of intermediate goods fixed. A low-income household devotes a fraction γ of each additional dollar of expenditure to mass-produced goods. A high-income household, by contrast, already consumes every existing variety at the bliss point and directs all additional expenditure toward luxury services. Redistributing resources toward the high-income type therefore reduces low-income demand for mass-produced goods without increasing the high-income type's purchases of each variety.

The production structure amplifies this shift in household demand. Greater demand for luxury services draws labor away from materials production and into the luxury sector. With less labor, materials producers demand fewer intermediate inputs at the prevailing price. Intermediate producers therefore face a contraction in price-sensitive demand from both households and materials firms. At the same time, an increase in luxury-sector dividends feeds back into household expenditure and partially offsets this contraction, but does not reverse it. In the Appendix we establish this result while fully accounting for the income feedback.

These shifts alter firms' pricing incentives. Because purchases by households at the bliss

point remain fixed, a larger share of each producer's sales becomes locally price inelastic. Raising the price therefore sacrifices fewer elastic sales relative to the fixed quantity purchased at bliss, making a higher markup optimal. This mechanism links the composition of demand to the common price charged to all buyers.

To determine the equilibrium response, we focus on locally stable equilibria: at a common price slightly above its equilibrium level, firms have an incentive to reduce their prices, whereas at a price slightly below equilibrium, they have an incentive to raise them. Along any regular branch of such equilibria, the contraction in elastic demand raises the equilibrium price:

$$\frac{dP}{ds} > 0. \tag{75}$$

The higher price restores firms' optimal pricing condition following the redistribution. The Appendix formalizes the stability criterion and states the regularity conditions underlying this local comparison. This criterion selects among balanced growth equilibria; it does not characterize the economy's transition from one equilibrium to another.

7.2.3 The Growth Effect

Combining the two steps yields the growth result. By reducing elastic demand relative to the fixed quantity purchased by households at the bliss point, redistribution raises the common price of intermediate goods. Although the markup increases, profits per variety decline, weakening the incentive to create new goods.

Proposition 3 (Inequality and innovation growth) *Along a regular, locally stable branch of interior BGPs, and holding population shares and all other primitives fixed, an increase in s—which redistributes income toward the high-income type—reduces the growth rate:*

$$\frac{dg}{ds} < 0. \tag{76}$$

The result follows from equation (72); The Appendix provides the formal proof. Equation (44) implies that the real interest rate declines. This comparison is local and applies as long as the mixed household regime, positive innovation, and global price optimality are preserved. High-income households continue to purchase newly introduced varieties, but concentrating income among them weakens the demand that sustains the profitability of each new product.

8 Calibration and Quantitative Implications

In this section, we study how greater income inequality affects demand, profits, the creation of new varieties, and growth. We hold the aggregate units of efficiency labor fixed, thereby isolating this distributional channel from changes in the economy's productivity.

The quantitative exercise connects the expenditure patterns documented in Section 4 to the model developed in Sections 5–7. Our microdata show that affluent households devote a larger share of their consumption expenditure to luxury goods and services, whose producers exhibit lower patent-based innovation intensity. We use the expenditure evidence to discipline the model's demand system and the patent evidence to motivate the distinction between innovation-intensive mass-produced goods and luxury services.

8.1 Mapping the Micro Evidence into the Calibration

Table 2 lists the empirical targets matched by the model, distributional and growth inputs, and maintained parameter values. We begin with the empirical evidence that disciplines household spending behavior.

Empirical discipline. The linked household–firm data described in Section 3 allow us to observe both the composition of household consumption and the innovation intensity of the firms receiving that expenditure. Two findings guide the calibration.

First, luxury expenditure shares are small and approximately constant throughout most of the distribution but rise sharply at the top. This pattern emerges both for firms on Deloitte's luxury list and under our broader classification, which identifies luxury firms by the concentration of their sales among affluent households (Figures 1 and 3). The model captures the flat portion through a common expenditure share below the bliss point and the increase at the top through satiation in existing mass-produced varieties.

Second, luxury firms exhibit lower patent intensity than other firms (Figures 2 and 4). Moreover, the consumption bundles of households at the top of the distribution are less innovation intensive according to citation-weighted, patent-value, and imputed-value measures (Figures 6, 7, and 8). These findings motivate the sectoral distinction through which the composition of expenditure affects innovation. The assumption that patented varieties are created only in the mass-produced sector is a stylized representation of this contrast: the evidence establishes that innovation intensity is lower in the luxury sector, not that every

luxury firm undertakes no innovation.

Household groups. We define the high-income type as households in the top five percent of the distribution ($m_H = 0.05$) and the low-income type as the remaining ninety-five percent ($m_L = 0.95$), matching the partition used to construct the empirical calibration moments.

The two expenditure-share moments. We calibrate the model to match the fractions of each group's consumption expenditure devoted to luxury:

$$\widehat{x}_{Q,L} = 0.013, \tag{77}$$

$$\widehat{x}_{Q,H} = 0.232. \tag{78}$$

The bottom 95 percent devote approximately 1.3 percent of their consumption budget to luxury goods and services, compared with 23.3 percent among the top 5 percent—a striking illustration of how Bernard Arnault, chairman and CEO of LVMH, describes luxury: "the ordinary for extraordinary people and the extraordinary for ordinary people."

These group-specific budget shares should be distinguished from the high-income group's share of aggregate income and from its share of total luxury sales. Their model counterparts are $x_{Q,j} = \frac{p_q q_j}{Z\epsilon_j}, \qquad j \in \{L, H\}$. We match these moments at the later calibration endpoint, which we label 2020.[6]

The two expenditure shares discipline different features of the demand system. We maintain that the low-income type remains below the bliss point at the later endpoint, so its luxury expenditure share is $1 - \gamma$. Matching the corresponding empirical moment therefore gives

$$1 - \gamma = \widehat{x}_{Q,L}, \qquad \gamma = 0.98666074. \tag{79}$$

Economically, this moment identifies the composition of expenditure before satiation alters the allocation of the consumption budget. It implies that households below the bliss point devote approximately 98.67 percent of their expenditure to mass-produced goods.

The much larger luxury expenditure share of the high-income group disciplines the location of the bliss point relative to that group's resources. Because preferences are identical, the two types would have the same expenditure shares if both remained below the bliss point. The shares would also coincide at the point at which the high-income type first reaches it. Conditional on the low-income type remaining below bliss, matching $\widehat{x}_{Q,H} > \widehat{x}_{Q,L}$ therefore

[6]The underlying microeconomic sample described in Section 3 covers 2015–2019.

requires the high-income type to lie strictly beyond the bliss point at the later endpoint. The cross-sectional difference in expenditure shares thus motivates the mixed regime used in the later calibration.

In this regime, the high-income type consumes each mass-produced variety at the bliss point, $\bar{c}$, so expenditure on these goods accounts for the fraction

$$\frac{P_{2020}\bar{c}}{Z_{2020}\epsilon_{H,2020}} = 1 - \widehat{x}_{Q,H} = 0.76701899 \tag{80}$$

of its consumption budget, compared with $\gamma = 0.98666074$ for the low-income type. Matching the high-income group's expenditure share therefore determines spending on mass-produced goods at the bliss point relative to that group's total expenditure. Equivalently,

$$\bar{c} = (1 - \widehat{x}_{Q,H})\frac{Z_{2020}s_{2020}}{m_H P_{2020}}. \tag{81}$$

The microdata do not directly reveal a physical satiation quantity for each variety. Instead, they identify the normalized expenditure object in (80), which the model maps into $\bar{c}$ through endogenous prices and expenditure. Because these equilibrium objects depend on research productivity, $\bar{c}$ and Ψ must be determined jointly. The calibration therefore uses the full endpoint equilibrium system rather than inferring the bliss point from the expenditure share alone.

Income distribution and the historical experiment. The expenditure moments describe how the two groups allocate their consumption budgets. Quantifying redistribution also requires specifying the share of the economy's resources received by each group. These are distinct inputs. We set $s_{1980} = 0.10$ and $s_{2020} = 0.17$. The population share of the high-income type remains fixed at $m_H = 0.05$, while its share of resources rises from 10 to 17 percent. At each value of s, the productivity levels satisfy

$$\epsilon_H(s) = \frac{s}{0.05}, \qquad \epsilon_L(s) = \frac{1-s}{0.95}. \tag{82}$$

The productivity of the high-income type rises from 2 to 3.4, while that of the low-income type falls from 0.9474 to 0.8737, leaving aggregate efficiency labor unchanged. The labels 1980 and 2020 refer to the two endpoint distributions; the calculations compare stationary allocations at these endpoints rather than tracing an annual transition between them. In particular, the later microeconomic sample identifies neither the 1980 income distribution nor the initial growth rate.

Technology and the growth target. The microeconomic expenditure moments discipline preferences but do not identify all production parameters or the return to creating a new variety. We take materials as the numeraire, normalize $\kappa = 1$, and set the annual rate of time preference to $\rho = 0.01$. We set $\nu = 1.4$, which implies a substitution elasticity of $\sigma = \nu/(\nu - 1) = 3.5$ and, by (21), $\alpha = 2/7$. When all buyers remain below the bliss point, the gross markup in the intermediate-goods sector is therefore ν. The assigned gross markup in luxury services, $\mu_Q = 1.525$, implies

$$\beta = 1 - \mu_Q^{-1} = 0.344262, \qquad \delta \equiv \frac{\beta}{1-\beta} = \mu_Q - 1 = 0.525. \tag{83}$$

These are maintained assumptions about technology and preferences, rather than parameters estimated from the microeconomic luxury-share moments or the patent-intensity comparisons.

We use the initial annual growth target, $g_{1980} = 0.02$, to pin down research productivity. Conditional on any candidate initial regime,

$$\Psi \pi_{x,1980} = \rho + g_{1980} = 0.03. \tag{84}$$

We impose no growth target at the later endpoint. Instead, the model predicts how much growth changes when the distribution of resources shifts, given a demand system disciplined by the microdata.

8.2 Determining the Initial Household Regime

We calibrate the economy separately under three candidate initial regimes: both household types below the bliss point (UU), the high-income type exactly at the bliss-point kink (KINK), and the high-income type strictly constrained by bliss while the low-income type remains below it (BU). Each candidate uses the same targets and maintained parameters but permits a different pair $(\Psi, \bar{c})$. At both endpoints, the calculation imposes luxury-market clearing and the pricing condition appropriate to the candidate regime. Under UU, $P_{1980} = \nu\kappa$. At KINK, the boundary condition for the high-income type replaces the smooth pricing equation, and the one-sided derivatives of profits must have the signs required for a local maximum.

The numerical procedure searches from multiple initial values and records all distinct roots it finds. For each root, it checks the positivity of quantities and expenditure, the

Table 2: Empirical moments, model counterparts, and calibration roles

Moment or parameter	Value	Mapping and role
Micro expenditure moments and household grouping		
High-type population share, m_H	0.05	Top-five-percent comparison
Low-type luxury share, $\widehat{x}_{Q,L}$	1.3339%	$1-\gamma$; directly identifies γ
High-type luxury share, $\widehat{x}_{Q,H}$	23.2981%	$1-P\bar{c}/(Z\epsilon_H)$; disciplines $\bar{c}$
Distributional inputs and initial growth		
Initial high-type resource share, s_{1980}	0.10	Initial distribution of resources
Later high-type resource share, s_{2020}	0.17	Later distribution of resources
Initial annual growth, g_{1980}	2.0000%	$\Psi\pi_{x,1980}-\rho$; disciplines Ψ
Assigned and implied parameters		
Discount rate, ρ	0.01	Assigned
Demand-curvature parameter, ν	1.40	Assigned
Gross luxury markup, μ_Q	1.525	Assigned
Intermediate materials requirement, κ	1	Normalization
Materials exponent, α	0.285714	$(\nu-1)/\nu$
Luxury-services parameter, β	0.344262	$1-\mu_Q^{-1}$
Ordinary-goods preference weight, γ	0.986661	$1-\widehat{x}_{Q,L}$
Jointly calibrated parameters		
Research productivity, Ψ	0.517131	Jointly calibrated
Bliss quantity per variety, $\bar{c}$	0.129481	Jointly calibrated

Notes: The expenditure shares are the direct microeconomic targets. The distributional inputs and the initial growth rate provide additional discipline. The patent-based comparisons motivate the model's sectoral mechanism but are not used as targets for Ψ. The reported parameter pair satisfies the implemented local pricing conditions. Quantity levels are expressed in model units. Neither later growth nor the initial household regime is targeted. Global price optimality imposes an additional requirement; see Section **??**.

Table 3: Calibration under alternative 1980 regimes

1980 regime	Ψ	$\bar{c}$	Result of implemented checks
UU	0.517131	0.129481	Passes the local-pricing checks.
KINK	0.546116	0.127656	Fails: the right-hand profit derivative is positive, $G_{UU,1980} = 0.03919$.
BU	0.499542	0.130691	Fails: $Q_{1980} = -0.01921$ and the high type's unconstrained demand is below bliss.

Notes: Each row reports a separate recalibration of the structural parameters. The multistart procedure finds one distinct numerical root for each candidate regime. Satisfying the reported checks does not establish global monopoly optimality at both endpoints.

household-regime inequalities, the relevant local pricing conditions, market-clearing and budget residuals, and the local stability of smooth BU allocations. The procedure continues only if exactly one candidate satisfies all these conditions. This approach distinguishes solutions to the equilibrium equations from those that also satisfy the implemented admissibility restrictions. It does not, however, prove that no additional roots exist outside the region covered by the numerical search.

Only the UU candidate satisfies the implemented checks. In this allocation, the high-income type's unconstrained consumption of each variety equals 79.25 percent of $\bar{c}$ in 1980 but would reach 128.64 percent of $\bar{c}$ in 2020. The latter is the symmetric unconstrained benchmark used to determine the household regime, not the quantity actually purchased once the bliss point binds. The low-income type remains below the bliss point at both endpoints.

The selected initial allocation therefore lies in the homothetic region of preferences, although the bliss point remains part of the preference structure throughout. Conditional on the selected UU regime, the initial growth target determines Ψ independently of the income distribution. Given Ψ, the later expenditure-share target then determines $\bar{c}$.

8.3 From the Initial to the Later Income Distribution

Holding the calibrated parameters fixed, we raise s from 0.10 to 0.17. At each value of s, the numerical procedure determines the relevant household regime and applies the corresponding local pricing condition. Table 4 reports the two endpoints and the boundaries between regimes.

Before the bliss point. As long as both household types remain below the bliss point, they allocate the same shares of their budgets to each consumption category. Proposition 2

Table 4: Stationary allocations along the local-pricing benchmark

	1980	First contact	Kink exit	2020
High-type resource share, s (%)	10.0000	12.6191	13.3601	17.0000
Regime	UU	UU/KINK	KINK/BU	BU
Common price, P	1.40000	1.40000	1.42861	1.42969
Luxury output, Q	0.01199	0.01199	0.01215	0.04680
Annual variety growth, g (%)	2.0000	2.0000	2.0092	1.9144
High luxury share (%)	1.3339	1.3339	1.3339	23.2981
Low luxury share (%)	1.3339	1.3339	1.3339	1.3339
Aggregate luxury share (%)	1.3339	1.3339	1.3339	5.0678
Relative innovation intensity	1.0000	1.0000	1.0000	0.7774

Notes: The first and last columns correspond to the endpoint income distributions. The intermediate columns report threshold values of s, not estimated calendar dates. Because luxury services are produced one-for-one with efficiency labor, luxury output equals employment measured in efficiency units; it need not equal the share of workers employed in the sector. Relative innovation intensity is defined in (87). The table reports the local-pricing branch generated by the numerical procedure; the 2020 allocation fails the global pricing check described in Section **??**.

therefore applies: redistribution leaves aggregate demand, the common price, and the growth rate of varieties unchanged. The high-income type first reaches the bliss point at

$$s_1 = \frac{m_H P_0 \bar{c}}{\gamma Z_0} = 0.126191, \tag{85}$$

where P_0 and Z_0 denote the price and normalized expenditure in the unconstrained stationary equilibrium. The initial resource share of 10 percent lies below this threshold.

The kink interval. Between first contact and $s_2 = 0.133601$, the high-income type remains exactly at the bliss-point boundary. Its consumption and expenditure satisfy

$$c_H = \bar{c}, \qquad P m_H \bar{c} = \gamma Z s. \tag{86}$$

As the high-income type receives more resources, the common price rises, keeping its expenditure share on mass-produced goods equal to γ. Consequently, both household types continue to allocate the share $1 - \gamma$ of their budgets to luxury services throughout this interval, even as prices and quantities adjust.

Over the kink interval, the price rises from 1.40000 to 1.42861, while growth increases slightly, from 2.0000 to 2.0092 percent. This increase does not contradict the comparative static for the smooth BU regime, which relies on an interior pricing equality that does not hold at the kink. Instead, a local optimum at the kink is characterized by one-sided derivative inequalities, which impose no general sign restriction on the growth response.

Prices, quantities, and growth are continuous at both boundaries, although their slopes need not be.

After the kink. Once the high-income type is strictly constrained by the bliss point, any further concentration of resources shifts its expenditure toward luxury services. Along the selected smooth local branch, the common price rises, while profits per variety decline. The local pricing identities derived in Section 7 therefore imply a negative growth response on this branch.

The later growth rate is 1.91436 percent, compared with the initial target of 2 percent—a decline of 0.08564 percentage point. The associated equilibrium interest rate falls from 3 to 2.91436 percent. Thus, along the local-pricing branch, the model accounts for approximately one-quarter of the observed decline in U.S. real per capita growth, from 2 percent in the 1980s to 1.65 percent in the 2010s. The decline in the equilibrium real interest rate reflects the deterioration in the economy's long-run growth prospects.

These results are consistent with evidence of diminishing dynamism in the U.S. economy before the recent AI boom. Gordon (2015) estimates that U.S. potential real GDP growth had fallen to approximately 1.6 percent, roughly half its average rate during 1972–2004. Byrne et al. (2016) document a slowdown in productivity growth, while Laubach and Williams (2003) argue that the equilibrium real interest rate has declined over time. Davis and Haltiwanger (2024) also document a decline in business formation rates.

The profit calculation makes the mechanism transparent. Between the two endpoints, total sales per intermediate producer fall from 0.14503 to 0.13115, while the price–cost margin rises from 0.40000 to 0.42969. Their product,

$$\pi_x = (P - \kappa)D,$$

declines from 0.05801 to 0.05636. The expansion of luxury production draws efficiency labor away from materials production, reducing materials producers' demand for intermediate inputs. Consumption per variety by the low-income type also falls. The higher margin partially offsets the decline in sales but is not sufficient to prevent profits from falling.

Luxury output rises from 0.01199 to 0.04680, increasing the sector's share of efficiency labor from 1.20 to 4.68 percent. Its share of consumption expenditure rises from 1.33 to 5.07 percent. These shares differ because prices and markups vary across sectors.

The model's connection to the empirical innovation Engel curve operates through changes

in the composition of household expenditure. Spending on mass-produced goods is associated with intermediate-goods patents, whereas luxury services generate no new patented varieties. Because all mass-produced varieties have the same patent intensity under symmetry, the relative innovation intensity of the two groups' expenditure bundles is

$$\frac{V_H}{V_L} = \frac{1 - x_{Q,H}}{1 - x_{Q,L}}. \tag{87}$$

This ratio declines from one to 0.77739, implying that the high-income type's expenditure bundle is approximately 22.26 percent less innovation intensive at the later endpoint. The model thus reproduces the direction of the expenditure-composition pattern documented in Figures 6, 7, and 8.

8.4 Higher Concentrations of Income

Holding the structural parameters fixed, the numerical procedure can extend the same local BU branch beyond $s = 0.17$. The continuation ends at the first implemented boundary, where variety growth reaches zero. This occurs at $s = 0.74453$. At that point, the common price is 1.53331 and luxury output is 0.75520. The high-income type allocates 83.00 percent of its expenditure to luxury services, while the aggregate luxury expenditure share is 62.14 percent.

9 Conclusions

In this paper, we show that the innovation intensity of the consumption bundle of the top 1 percent is substantially lower than that of the remaining households. Through the lens of the model, this finding implies that redistribution toward the very top of the income distribution changes not only the level of demand but also its composition: it shifts expenditure away from goods that are relatively innovation-intensive and toward goods for which marginal demand provides weaker incentives for innovation. In this sense, inequality can affect growth through the direction of demand.

This mechanism has broader implications for how we think about the relationship between inequality and growth. We discuss some of these implications informally here, leaving a more complete analysis for future research.

First, population size and income per capita are not interchangeable when aggregate income is held fixed. In the model, reducing population while increasing income per capita

can push a larger fraction of households beyond the income thresholds at which marginal spending shifts away from innovation-intensive goods. A smaller but richer population may therefore generate less demand for the types of products that sustain broad-based innovation and growth.

Second, the Gini coefficient is not a sufficiently informative measure of inequality in this economy. What matters for innovation is not only the overall dispersion of income, but also where income is located in the distribution. It makes a fundamental difference whether 15 percent of total income accrues to the top 1 percent or the top 10 percent, even if the Gini coefficient is similar in both cases. The reason is that households at different points in the distribution may have very different marginal propensities to spend on innovation-intensive goods.

Third, the mechanism by which inequality affects growth remains relevant in an open economy. One might argue that, if firms can sell globally, only global demand should matter for innovation. However, in practice, exporters often develop, test, and refine products in their domestic markets before succeeding abroad (see, e.g., Bernard and Jensen (1999)). Domestic demand can therefore shape the direction of innovation even in economies that are highly integrated into world markets.

Finally, the model offers a perspective on why high inequality can be costly in developing economies. When income is highly concentrated, demand for mass-consumption goods is weaker, reducing incentives to innovate. Exporting can mitigate this force, but may not fully offset it if firms must first build capabilities at home before competing internationally. This perspective is consistent with the experience of several rapidly growing economies, including the Asian Tigers (Hong Kong, Singapore, South Korea, and Taiwan), which combined export orientation with relatively low levels of income inequality during important phases of their development.

References

Acemoglu, D. and V. Guerrieri (2008): "Capital Deepening and Nonbalanced Economic Growth," *Journal of Political Economy*, 116, 467–498.

Aghion, P., E. Caroli, and C. García-Peñalosa (1999): "Inequality and economic growth: the perspective of the new growth theories," *Journal of Economic Literature*, 37, 1615–1660.

Aglietta, M. (1979): *A Theory of Capitalist Regulation: The US Experience*, Verso, originally published in French in 1976.

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

Arora, S. and M. Dell (2023): "Linking Records with Natural Language: A Transformer-Based Approach," ArXiv preprint arXiv:2309.00789.

Bernard, A. B. and J. B. Jensen (1999): "Exceptional Exporter Performance: Cause, Effect, or Both?" *Journal of International Economics*, 47, 1–25.

Berry, C. J. (1994): *The Idea of Luxury: A Conceptual and Historical Investigation*, Cambridge: Cambridge University Press.

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

Buera, F. J. and J. P. Kaboski (2012): "The Rise of the Service Economy," *American Economic Review*, 102, 2540–2569.

Byrne, D. M., J. G. Fernald, and M. B. Reinsdorf (2016): "Does the United States Have a Productivity Slowdown or a Measurement Problem?" *Brookings Papers on Economic Activity*, 2016, 109–182.

Christiano, L. J., M. Eichenbaum, and D. Marshall (1991): "The Permanent Income Hypothesis Revisited," *Econometrica: Journal of the Econometric Society*, 397–423.

Cochrane, J. H. (2014): "A mean-variance benchmark for intertemporal portfolio theory," *the Journal of Finance*, 69, 1–49.

Davis, S. J. and J. Haltiwanger (2024): "Dynamism Diminished: The Role of Housing Markets and Credit Conditions," *American Economic Journal: Macroeconomics*, 16, 29–61.

Deaton, A. and J. Muellbauer (1980): *Economics and Consumer Behavior*, Cambridge University Press.

Duarte, M. and D. Restuccia (2010): "The Role of the Structural Transformation in Aggregate Productivity," *Quarterly Journal of Economics*, 125, 129–173.

Engel, E. (1857): *Die Productions- und Consumtionsverhältnisse des Königreichs Sachsen*, C. Heinrich.

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

Foellmi, R. and J. Zweimüller (2008): "Structural change, Engel's consumption cycles and Kaldor's facts of economic growth," *Journal of monetary Economics*, 55, 1317–1328.

Gollin, D., S. Parente, and R. Rogerson (2002): "The role of agriculture in development," *American economic review*, 92, 160–164.

Gollin, D., S. L. Parente, and R. Rogerson (2007): "The food problem and the evolution of international income levels," *Journal of Monetary Economics*, 54, 1230–1255.

Gordon, R. J. (2015): "Secular Stagnation: A Supply-Side View," *American Economic Review*, 105, 54–59.

——— (2016): *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War*, Princeton University Press.

Hall, R. E. (1978): "Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence," *Journal of political economy*, 86, 971–987.

Hallak, J. C. (2006): "Product quality and the direction of trade," *Journal of international Economics*, 68, 238–265.

Hayek, F. A. (1960): *The Constitution of Liberty*, Chicago: University of Chicago Press.

Heaton, J. (1993): "The interaction between time-nonseparable preferences and time aggregation," *Econometrica: Journal of the Econometric Society*, 353–385.

Howard, J. A. and J. N. Sheth (1969): *The Theory of Buyer Behavior*, John Wiley & Sons.

Hume, D. (1987a): "Of Commerce," in *Essays: Moral, Political, and Literary*, ed. by E. F. Miller, Indianapolis: Liberty Fund, originally published in Political Discourses, 1752.

——— (1987b): "Of Refinement in the Arts," in *Essays: Moral, Political, and Literary*, ed. by E. F. Miller, Indianapolis: Liberty Fund, originally published in Political Discourses, 1752; earlier versions titled "Of Luxury".

Imas, A. and K. Madarász (2024): "Superiority-Seeking and the Preference for Exclusion," *The Review of Economic Studies*, 91, 2347–2382.

Jaimovich, N., S. Rebelo, A. Wong, and M. Zhang (2019): "Trading Up and the Skill Premium," *NBER Macro Annual*.

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

Kaldor, N. (1966): "Causes of the Slow Rate of Economic Growth of the United Kingdom," in *Causes of the Slow Rate of Economic Growth of the United Kingdom*, Cambridge University Press.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017): "Technological Innovation, Resource Allocation, and Growth," *The Quarterly Journal of Economics*, 132, 665–712.

Kongsamut, P., S. Rebelo, and D. Xie (2001): "Beyond Balanced Growth," *Review of Economic Studies*, 68, 869–882.

Kotler, P. (1967): *Marketing Management: Analysis, Planning, and Control*, Prentice-Hall.

Kuznets, S. (1955): "Economic Growth and Income Inequality," *American Economic Review*, 45, 1–28.

Laitner, J. (2000): "Structural change and economic growth," *The Review of Economic Studies*, 67, 545–561.

Laubach, T. and J. C. Williams (2003): "Measuring the Natural Rate of Interest," *The Review of Economics and Statistics*, 85, 1063–1070.

Mandel, G. and A. Imas (2026): "Art and the Machine: Why People Devalue AI-Generated Creative Work," Manuscript.

Mankiw, N. G. (1982): "Hall's consumption hypothesis and durable goods," *Journal of Monetary Economics*, 10, 417–425.

Marshall, M. G. (2000): "Luxury, Economic Development, and Work Motivation: David Hume, Adam Smith, and J. R. McCulloch," *History of Political Economy*, 32, 631–648.

Maslow, A. H. (1943): "A Theory of Human Motivation," *Psychological Review*, 50, 370–396.

Ngai, L. R. and C. A. Pissarides (2007): "Structural Change in a Multisector Model of Growth," *American Economic Review*, 97, 429–443.

Piketty, T. and E. Saez (2003): "Income Inequality in the United States, 1913–1998," *Quarterly Journal of Economics*, 118, 1–39.

Pollak, R. A. (1971): "Additive Utility Functions and Linear Engel Curves," *Review of Economic Studies*, 38, 401–414.

Ramsey, F. P. (1928): "A mathematical theory of saving," *The Economic Journal*, 38, 543–559.

Romer, P. M. (1990): "Endogenous Technological Change," *Journal of Political Economy*, 98, S71–S102.

Samuelson, P. A. and R. M. Solow (1956): "A complete capital model involving heterogeneous capital goods," *The Quarterly Journal of Economics*, 70, 537–562.

Schumpeter, J. A. (1950): *Capitalism, Socialism and Democracy*, New York: Harper & Brothers, 3 ed.

Smith, A. (1982): *The Theory of Moral Sentiments*, Indianapolis: Liberty Fund, originally published in 1759.

Verboven, F. (1996): "International price discrimination in the European car market," *The RAND Journal of Economics*, 240–268.